

Logon

*** It is now 12/12/08 10:12:46 AM ***

Welcome to DialogLink - Version 5 Revolutionize the Way You Work!

New on Dialog

Order Patent and Trademark File Histories Through Dialog

Thomson File Histories are now available directly through *Dialog*. Combined with the comprehensive patent and trademark information on *Dialog*, file histories give you the most complete view of a patent or trademark and its history in one place. When searching in the following patent and trademark databases, a link to an online order form is displayed in your search results, saving you time in obtaining the file histories you need.

Thomson File Histories are available from the following *Dialog* databases:

- CLAIMS/Current Patent Legal Status (File 123)
- CLAIMS/U.S. Patents (File 340)
- Chinese Patent Abstracts in English (File 344)
- Derwent Patents Citation Index (File 342)
- Derwent World Patents Index (for users in Japan) (File 352)
- Derwent World Patents Index First View (File 331)
- Derwent World Patents Index (File 351)
- Derwent World Patents Index (File 350)
- Ei EnCompassPat (File 353)
- European Patents Fulltext (File 348)
- French Patents (File 371)
- German Patents Fulltext (File 324)
- IMS Patent Focus (File 447, 947)
- INPADOC/Family and Legal Status (File 345)
- JAPIO - Patent Abstracts of Japan (File 347)
- LitAlert (File 670)
- U.S. Patents Fulltext (1971-1975) (File 652)

- U.S. Patents Fulltext (1976-present) (File 654)
- WIPO/PCT Patents Fulltext (File 349)
- TRADEMARKSCAN - U.S. Federal (File 226)

DialogLink 5 Release Notes

New features available in the latest release of DialogLink 5 (August 2006)

- Ability to resize images for easier incorporation into DialogLink Reports
- New settings allow users to be prompted to save Dialog search sessions in the format of their choice (Microsoft Word, RTF, PDF, HTML, or TEXT)
- Ability to set up Dialog Alerts by Chemical Structures and the addition of Index Chemicus as a structure searchable database
- Support for connections to STN Germany and STN Japan services

Show Preferences for details

? Help Off Line

* * *

Connecting to Rob Pond - Dialog - 264751

Connected to Dialog via SMS002043286

? B 15, 9, 610, 810, 275, 476, 624, 621, 636, 613, 813, 16, 160, 634, 148, 20, 35, 583, 65, 2,474, 475, 99, 256, 635, 570, PAPERSMJ, PAPERSEU, 47

>>>W: 476 does not exist

1 of the specified files is not available

[File 15] ABI/Inform(R) 1971-2008/Dec 10

(c) 2008 ProQuest Info&Learning. All rights reserved.

[File 9] Business & Industry(R) Jul/1994-2008/Dec 11

(c) 2008 Gale/Cengage. All rights reserved.

[File 610] Business Wire 1999-2008/Dec 12

(c) 2008 Business Wire. All rights reserved.

**File 610: File 610 now contains data from 3/99 forward. Archive data (1986-2/99) is available in File 810.*

[File 810] Business Wire 1986-1999/Feb 28

(c) 1999 Business Wire . All rights reserved.

[File 275] Gale Group Computer DB(TM) 1983-2008/Nov 26
(c) 2008 Gale/Cengage. All rights reserved.

[File 624] McGraw-Hill Publications 1985-2008/Dec 11
(c) 2008 McGraw-Hill Co. Inc. All rights reserved.

[File 621] Gale Group New Prod. Annou.(R) 1985-2008/Nov 14
(c) 2008 Gale/Cengage. All rights reserved.

[File 636] Gale Group Newsletter DB(TM) 1987-2008/Dec 01
(c) 2008 Gale/Cengage. All rights reserved.

[File 613] PR Newswire 1999-2008/Dec 12
(c) 2008 PR Newswire Association Inc. All rights reserved.

**File 613: File 613 now contains data from 5/99 forward. Archive data (1987-4/99) is available in File 813.*

[File 813] PR Newswire 1987-1999/Apr 30
(c) 1999 PR Newswire Association Inc. All rights reserved.

[File 16] Gale Group PROMT(R) 1990-2008/Dec 01
(c) 2008 Gale/Cengage. All rights reserved.
**File 16: Because of updating irregularities, the banner and the update (UD=) may vary.*

[File 160] Gale Group PROMT(R) 1972-1989
(c) 1999 The Gale Group. All rights reserved.

[File 634] San Jose Mercury Jun 1985-2008/Dec 08
(c) 2008 San Jose Mercury News. All rights reserved.

[File 148] Gale Group Trade & Industry DB 1976-2008/Dec 08
(c) 2008 Gale/Cengage. All rights reserved.
**File 148: The CURRENT feature is not working in File 148. See HELP NEWS148.*

[File 20] Dialog Global Reporter 1997-2008/Dec 12
(c) 2008 Dialog. All rights reserved.

[File 35] Dissertation Abs Online 1861-2008/Feb
(c) 2008 ProQuest Info&Learning. All rights reserved.

[File 583] Gale Group Globalbase(TM) 1986-2002/Dec 13
(c) 2002 Gale/Cengage. All rights reserved.
**File 583: This file is no longer updating as of 12-13-2002.*

[File 65] Inside Conferences 1993-2008/Dec 11
(c) 2008 BLDSC all rts. reserv. All rights reserved.

[File 2] INSPEC 1898-2008/Nov W3
(c) 2008 Institution of Electrical Engineers. All rights reserved.

[File 474] New York Times Abs 1969-2008/Dec 12
(c) 2008 The New York Times. All rights reserved.

[File 475] Wall Street Journal Abs 1973-2008/Dec 12
(c) 2008 The New York Times. All rights reserved.

[File 99] Wilson Appl. Sci & Tech Abs 1983-2008/Oct
(c) 2008 The HW Wilson Co. All rights reserved.

[File 256] TecInfoSource 82-2008/Jul

(c) 2008 Info.Sources Inc. All rights reserved.

[File 635] Business Dateline(R) 1985-2008/Dec 11

(c) 2008 ProQuest Info&Learning. All rights reserved.

[File 570] Gale Group MARS(R) 1984-2008/Dec 01

(c) 2008 Gale/Cengage. All rights reserved.

[File 387] The Denver Post 1994-2008/Dec 10

(c) 2008 Denver Post. All rights reserved.

[File 471] New York Times Fulltext 1980-2008/Dec 12

(c) 2008 The New York Times. All rights reserved.

[File 492] Arizona Repub/Phoenix Gaz 19862002/Jan 06

(c) 2002 Phoenix Newspapers. All rights reserved.

**File 492: File 492 is closed (no longer updating). Use Newsroom, Files 989 and 990, for current records.*

[File 494] St LouisPost-Dispatch 1988-2008/Dec 11

(c) 2008 St Louis Post-Dispatch. All rights reserved.

[File 631] Boston Globe 1980-2008/Dec 11

(c) 2008 Boston Globe. All rights reserved.

[File 633] Phil.Inquirer 1983-2008/Dec 12

(c) 2008 Philadelphia Newspapers Inc. All rights reserved.

[File 638] Newsday/New York Newsday 1987-2008/Dec 11

(c) 2008 Newsday Inc. All rights reserved.

[File 640] San Francisco Chronicle 1988-2008/Dec 10

(c) 2008 Chronicle Publ. Co. All rights reserved.

[File 641] Rocky Mountain News Jun 1989-2008/Dec 12

(c) 2008 Scripps Howard News. All rights reserved.

[File 702] Miami Herald 1983-2008/Dec 12

(c) 2008 The Miami Herald Publishing Co. All rights reserved.

[File 703] USA Today 1989-2008/Dec 11

(c) 2008 USA Today. All rights reserved.

[File 704] (Portland)The Oregonian 1989-2008/Dec 10

(c) 2008 The Oregonian. All rights reserved.

[File 713] Atlanta J/Const. 1989-2008/Nov 09

(c) 2008 Atlanta Newspapers. All rights reserved.

[File 714] (Baltimore) The Sun 1990-2008/Dec 11

(c) 2008 Baltimore Sun. All rights reserved.

[File 715] Christian Sci.Mon. 1989-2008/Dec 10

(c) 2008 Christian Science Monitor. All rights reserved.

[File 725] (Cleveland)Plain Dealer Aug 1991-2008/Dec 11

(c) 2008 The Plain Dealer. All rights reserved.

[File 735] St. Petersburg Times 1989- 2008/Dec 07
(c) 2008 St. Petersburg Times. All rights reserved.

[File 477] Irish Times 1999-2008/Dec 11
(c) 2008 Irish Times. All rights reserved.

[File 710] Times/Sun.Times(London) Jun 1988-2008/Dec 09
(c) 2008 Times Newspapers. All rights reserved.

[File 711] Independent(London) Sep 1988-2006/Dec 12
(c) 2006 Newspaper Publ. PLC. All rights reserved.

**File 711: This file does not update. See File 757 for full daily coverage from many European sources.*

[File 756] Daily/Sunday Telegraph 2000-2008/Dec 12
(c) 2008 Telegraph Group. All rights reserved.

[File 757] Mirror Publications/Independent Newspapers 2000-2008/Dec 11
(c) 2008. All rights reserved.

[File 47] Gale Group Magazine DB(TM) 1959-2008/Dec 10
(c) 2008 Gale/Cengage. All rights reserved.

**File 47: UD names have been adjusted to reflect process dates All data is present*

? S pd<20000329

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

>>>W: One or more prefixes are unsupported

or undefined in one or more files.

S1 63914638 S PD<20000329

? S s1 and (BILL??? OR INVOIC??? OR CHARG??? OR PAYMENT OR PAYMENTS OR SETTL??? OR SETTLEMENT) AND (BROKER??? OR SYNCHRO OR SYNCHRONIZ??? OR SYNCHRONIZATION OR MEDIAT??? OR MEDIATION OR INTERMEDIAT???) AND (SERVER OR COMPUTER OR SERVERS OR COMPUTERIZ??? OR COMPUTERIZATION OR COMPUTERS OR PROCESS??? OR TERMINAL OR TERMINALS OR UNIT OR UNITS OR APPARATUS)

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

Processing

63914638 S1

22970346 BILL???

318853 INVOIC???

13438911 CHARG???

3363395 PAYMENT

2997399 PAYMENTS

3165580 SETTL???

2038739 SETTLEMENT

5947611	BROKER???
7021	SYNCHRO
260834	SYNCHRONIZ???
197278	SYNCHRONIZATION
702219	MEDIAT???
254767	MEDIATION
2247134	INTERMEDIAT???
2404233	SERVER
14201129	COMPUTER
1457674	SERVERS
427882	COMPUTERIZ???
37241	COMPUTERIZATION
7441971	COMPUTERS
22614619	PROCESS???
1412911	TERMINAL
838286	TERMINALS
8156498	UNIT
5465091	UNITS
1056949	APPARATUS

S2 361599 S S1 AND (BILL???

OR INVOIC???

OR CHARG???

OR PAYMENT OR PAYMENTS OR

SETTL???

OR SETTLEMENT) AND (BROKER???

OR SYNCHRO OR SYNCHRONIZ???

OR SYNCHRONIZATION OR

MEDIAT???

OR MEDIATION OR INTERMEDIAT???) AND (SERVER OR COMPUTER OR SERVERS OR

COMPUTERIZ???

OR COMPUTERIZATION OR COMPUTERS OR PROCESS???

OR TERMINAL OR TERMINALS OR

UNIT OR UNITS OR APPARATUS)

? s s2 and ((synchro or synchroniz???

or synchronization) (5n) (memory or buffer???)

Processing

361599	S2
7021	SYNCHRO
260834	SYNCHRONIZ???
197278	SYNCHRONIZATION
2655609	MEMORY
421896	BUFFER???
5563	((SYNCHRO OR SYNCHRONIZ???) OR SYNCHRONIZATION) (5N) (MEMORY OR BUFFER???)

S3 250 S S2 AND ((SYNCHRO OR SYNCHRONIZ??? OR SYNCHRONIZATION) (5N) (MEMORY OR BUFFER???))

? rd

S4 155 RD (UNIQUE ITEMS)

? t s4/free/all

4/8/1 (Item 1 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

02332051 86066948

USE FORMAT 7 OR 9 FOR FULL TEXT

Bottleneck allocation methodology (BAM): an algorithm

Word Count: 4251

1999

Geographic Names: United States; US

Descriptors: Algorithms; Studies; Manufacturing resource planning; Production scheduling

Classification Codes: 9190 (CN=United States); 9130 (CN=Experimental/Theoretical); 5310 (CN=Production planning & control)

Print Media ID: 11839

4/8/2 (Item 2 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01993204 50586330

USE FORMAT 7 OR 9 FOR FULL TEXT

Cell phones strike back

Word Count: 561 Length: 2 Pages

Feb 28, 2000

Company Names:

Intel Corp (Duns: 04-789-7855) (Ticker: INTC SIC:3674) (NAICS: 334413) (NAICS:334210) (

NAICS:334419) (NAICS:334611)

Geographic Names: United States; US

Descriptors: Cellular telephones; Handheld computers; Wireless communications; Communications equipment; Trends

Classification Codes: 5250 (CN=Telecommunications systems & Internet communications); 8650 (CN=Electrical & electronics industries); 9190 (CN=United States)

Print Media ID: 17765

4/8/3 (Item 3 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01887959 05-38951

USE FORMAT 7 OR 9 FOR FULL TEXT

SDRAM memory: DRAM and beyond

Word Count: 1541 Length: 4 Pages

Second Quarter 1999

Geographic Names: US

Descriptors: DRAM; R&D; Computer industry; Capacity; Bandwidths; Performance evaluation; Comparative analysis; Technological change

Classification Codes: 9190 (CN=United States); 5230 (CN=Computer hardware); 5400 (CN=Research & development); 8651 (CN=Computer industry)

4/8/4 (Item 4 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01838320 04-89311

USE FORMAT 7 OR 9 FOR FULL TEXT

Data capture grows wider

Word Count: 2241 Length: 5 Pages

Jun 14, 1999

Company Names:

Federal Express Corp (Duns: 05-807-0459 Ticker: FDX)

Beth Israel Deaconess Medical Center-Boston MA

Cummins Engine Co Inc (Duns: 00-641-5160 Ticker: CUM)

Brooks Brothers

Microsoft Corp (Duns: 08-146-6849 Ticker: MSFT)

Geographic Names: US

Descriptors: Data mining; Trends; Data warehouses; Portable computers; Manycompanies

Classification Codes: 9190 (CN=United States); 5240 (CN=Software & systems); 5220 (CN=Data processing management)

4/8/5 (Item 5 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01835242 04-86233

USE FORMAT 7 OR 9 FOR FULL TEXT

All-optical networks

Word Count: 5031 Length: 10 Pages

Jun 1999

Geographic Names: US

Descriptors: Fiber optic networks; Network topologies; Communications equipment ; Multiplexers; Data transmission

Classification Codes: 5250 (CN=Telecommunications systems); 9190 (CN=United States)

4/8/6 (Item 6 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01778495 04-29486

****USE FORMAT 7 OR 9 FOR FULL TEXT****

Always on

Word Count: 866 Length: 2 Pages

Feb 15, 1999

Company Names:

Nortel Networks

Motorola Computer Group

Geographic Names: US

Descriptors: Carriers ; Technological planning; Communications networks; Reliability

Classification Codes: 9190 (CN=United States); 8330 (CN=Broadcasting & telecommunications); 5250 (CN=Telecommunications systems); 2400 (CN=Public relations)

4/8/7 (Item 7 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01696282 03-47272

****USE FORMAT 7 OR 9 FOR FULL TEXT****

A pattern system for network management interfaces

Word Count: 4553 Length: 8 Pages

Sep 1998

Geographic Names: US

Descriptors: Network management systems; Application programming interface; Systems design; Studies

Classification Codes: 9190 (CN=United States); 5240 (CN=Software & systems); 9130 (CN=Experimental/Theoretical)

4/8/8 (Item 8 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01673556 03-24546

****USE FORMAT 7 OR 9 FOR FULL TEXT****

Java for all platforms

Word Count: 1384 Length: 3 Pages

Jul 20, 1998

Company Names:

Sun Microsystems Inc (Duns: 01-304-4532 Ticker: SUNW)

Geographic Names: US

Descriptors: Java; Systems portability; Technological change; Object oriented programming
Classification Codes: 9190 (CN=United States); 5240 (CN=Software & systems)

4/8/9 (Item 9 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01608568 02-59557

****USE FORMAT 7 OR 9 FOR FULL TEXT****

Maximizing markets: Quicker speed? Get faster memory

Word Count: 936 Length: 2 Pages

Mar 30, 1998

Geographic Names: US

Descriptors: DRAM; Industrywide conditions; Computer industry; Profits; Processing speed; Product development

Classification Codes: 8651 (CN=Computer industry); 9190 (CN=United States); 7500 (CN=Product planning & development)

4/8/10 (Item 10 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01514315 01-65303

****USE FORMAT 7 OR 9 FOR FULL TEXT****

SLDRAM Consortium puts up fight for memory

Word Count: 543 Length: 1 Pages

Oct 6, 1997

Geographic Names: US

Descriptors: Consortia; Standards; Computer memory; Computer architecture; Competition

Classification Codes: 5230 (CN=Computer hardware); 7500 (CN=Product planning & development); 9190 (CN=United States)

4/8/11 (Item 11 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01498286 01-49274

****USE FORMAT 7 OR 9 FOR FULL TEXT****

Sharing memory with memory-mapped files

Word Count: 2529 Length: 6 Pages

Oct 1997

Geographic Names: US

Descriptors: Distributed processing ; UNIX; Monte Carlo simulation; Computer programming; Problem solving;

Methods; Systems development

Classification Codes: 9190 (CN=United States); 5240 (CN=Software & systems)

4/8/12 (Item 12 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01487026 01-38014

****USE FORMAT 7 OR 9 FOR FULL TEXT****

Nortel's 10-GBIT/S transport platform: Delivering bandwidth to build on

Word Count: 4364 Length: 11 Pages

Jul 1997

Company Names:

Nortel Communications Inc

Geographic Names: Canada

Descriptors: Bandwidths; Multiplexers; Technological change; SONET; Product development

Classification Codes: 5250 (CN=Telecommunications systems); 8650 (CN=Electrical & electronics industries);

7500 (CN=Product planning & development); 5400 (CN=Research & development); 9172 (CN=Canada)

4/8/13 (Item 13 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01218374 98-67769

****USE FORMAT 7 OR 9 FOR FULL TEXT****

CMG 1995 annual conference reflects the state of the industry

Word Count: 5151 Length: 12 Pages

Jan 1996

Geographic Names: US

Descriptors: Systems management; Computer memory; Performance evaluation; Conferences

Classification Codes: 5240 (CN=Software & systems); 7300 (CN=Sales & selling); 9190 (CN=United States)

4/8/14 (Item 14 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01089259 97-38653

****USE FORMAT 7 OR 9 FOR FULL TEXT****

Testing wireless

Word Count: 2928 Length: 5 Pages

Sep 1995

Geographic Names: US

Descriptors: Wireless communications; Infrastructure ; Equipment testing; Maintenance management

Classification Codes: 9190 (CN=United States); 5250 (CN=Telecommunications systems); 5130

(CN=Maintenance)

4/8/15 (Item 15 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

00975539 96-24932

****USE FORMAT 7 OR 9 FOR FULL TEXT****

JIT's impact on a firm's financial statements

Word Count: 3312 Length: 5 Pages

Winter 1995

Descriptors: Studies; Purchasing; Just in time; Inventory management

Classification Codes: 5120 (CN=Purchasing); 9130 (CN=Experimental/Theoretical); 5330 (CN=Inventory management)

4/8/16 (Item 16 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

00901316 95-50708

****USE FORMAT 7 OR 9 FOR FULL TEXT****

Lead-time models of business processes

Word Count: 6133 Length: 16 Pages

1994

Descriptors: Operations research; Production planning; Time management; Business process reengineering; Management styles; Models

Classification Codes: 2600 (CN=Management science/Operations research); 5310 (CN=Production planning & control); 2200 (CN=Managerial skills); 9130 (CN=Experimental/Theoretical)

4/8/17 (Item 17 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

00842805 94-92197

****USE FORMAT 7 OR 9 FOR FULL TEXT****

Advances in parallel computing for reactor analysis and safety

Word Count: 6093 Length: 11 Pages

Apr 1994

Geographic Names: US

Descriptors: Systems development; Parallel processing ; Computer based modeling; Simulation; Nuclear reactors; Safety management; Applications

Classification Codes: 9190 (CN=United States); 5240 (CN=Software & systems); 8340 (CN=Electric, water & gas utilities); 5340 (CN=Safety management); 9130 (CN=Experimental/Theoretical)

4/8/18 (Item 18 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

00627563 92-42503

****USE FORMAT 7 OR 9 FOR FULL TEXT****

Ultracomputers: A Teraflop Before Its Time

Word Count: 13520 Length: 22 Pages

Aug 1992

Descriptors: R&D; Supercomputers; Computer industry; Product development; Parallel processing ; Processing speed; Multiprocessing

Classification Codes: 5400 (CN=Research & development); 5230 (CN=Computer hardware); 8651 (CN=Computer industry); 7500 (CN=Product planning & development)

4/8/19 (Item 19 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

00623706 92-38808

****USE FORMAT 7 OR 9 FOR FULL TEXT****

Vendors Reply to Frame Relay 20 Questions

Word Count: 936 Length: 2 Pages

Jul 1992

Company Names:

Cascade Communications Corp

Cisco Systems Inc (Duns: 15-380-4570 Ticker: CSCO)

Motorola-Codex

Netrix Corp

Sync Research

Geographic Names: US

Descriptors: Packet switched networks; Manyproducts; Manycompanies; Connectivity; Data transmission; Standards; Support

Classification Codes: 5250 (CN=Telecommunications systems); 9190 (CN=United States)

4/8/20 (Item 1 from file: 9)

Business & Industry(R)

(c) 2008 Gale/Cengage. All rights reserved.

01865006 Supplier Number: 24680837 (USE FORMAT 7 OR 9 FOR FULLTEXT)

Data Capture Grows Wider -- Small Computing Devices And Embedded Systems Can Feed Large Data Warehouses, Leading To Potentially Powerful Data Analysis

June 14, 1999

Word Count: 2141

Industry Names: Applications software; Computer; Mobile communications; Personal computers; Portable computers; Software; Telecom services; Telecommunications

Product Names: Portable computers (357165); Radiotelephone communications (481200); Database software packages (737265); Applications software packages NEC (737279)

Concept Terms: All market information; Industry forecasts; Sales; Trends; Users

Geographic Names: World (WOR)

4/8/21 (Item 2 from file: 9)

Business & Industry(R)

(c) 2008 Gale/Cengage. All rights reserved.

01593920 Supplier Number: 24317129 (USE FORMAT 7 OR 9 FOR FULLTEXT)

Silver lining seen in DRAM storm cloud

July 06, 1998

Word Count: 1401

Special Features: Table

Industry Names: Electronic components; Semiconductors

Product Names: Memory integrated circuits (367445)

Concept Terms: All market information; Industry forecasts; Market size; Sales

Geographic Names: World (WOR)

4/8/22 (Item 3 from file: 9)

Business & Industry(R)

(c) 2008 Gale/Cengage. All rights reserved.

00780751 Supplier Number: 23328070 (USE FORMAT 7 OR 9 FOR FULLTEXT)

Taligent Keeps Its Promises

October 23, 1995

Word Count: 2979

Company Names: APPLE COMPUTER INC; HEWLETT-PACKARD LTD (HEWLETT-PACKARD CO); INTERNATIONAL BUSINESS MACHINES CORP; TALIGENT INC

Industry Names: Applications software; Software

Product Names: Applications software packages (737263)

Concept Terms: All product and service information; Product introduction

Geographic Names: North America (NOAX); United States (USA)

4/8/23 (Item 1 from file: 610)

Business Wire

(c) 2008 Business Wire. All rights reserved.

00209224 20000306066B5505 (USE FORMAT 7 FOR FULLTEXT)

Synchrologic Announces First Total Systems Management Solution for Handhelds and Mobile PCs

Monday , March 6, 2000 16:11 EST

Word Count: 632

Company Names: PALM COMPUTING INC; US ROBOTICS CORP

Geographic Names: CALIFORNIA; AMERICAS; NORTH AMERICA; USA

Product Names: COMPUTER SOFTWARE; PORTABLE COMPUTERS; COMPUTERS; COMPUTER
HARDWARE; MICROCOMPUTERS

Event Names: TECHNOLOGY DEVELOPMENT

4/8/24 (Item 2 from file: 610)

Business Wire

(c) 2008 Business Wire. All rights reserved.

00206546 20000301061B2644 (USE FORMAT 7 FOR FULLTEXT)

WideBand Corp. Begins Trading

Wednesday , March 1, 2000 18:53 EST

Word Count: 407

Product Names: CORPORATE NETWORKS; NETWORKS; COMMUNICATIONS TECHNOLOGIES;
COMPUTERS ; CORPORATE; DATA COMMUNICATIONS

Event Names: TECHNOLOGY DEVELOPMENT

4/8/25 (Item 3 from file: 610)

Business Wire

(c) 2008 Business Wire. All rights reserved.

00107707 19990922265B0166 (USE FORMAT 7 FOR FULLTEXT)

Ericsson Introduces State-of-the-Art Phone for International Business Travelers

Wednesday , September 22, 1999 09:20 EDT

Word Count: 767

Company Names: TELEFON AB LM ERICSSON; OMNIPOINT CORP; EDELMAN PUBLIC RELATIONS

Geographic Names: USA; AMERICAS; NORTH AMERICA

Product Names: ELECTRONIC MAIL; MOBILE COMMUNICATIONS; NETWORKS; PORTABLE
COMPUTERS ; RADIO COMMUNICATION; TELEPHONES; COMMUNICATIONS TECHNOLOGIES;
COMPUTERS; DATA COMMUNICATIONS; TELECOMMUNICATIONS; COMPUTER HARDWARE;
MICROCOMPUTERS

4/8/26 (Item 4 from file: 610)

Business Wire

(c) 2008 Business Wire. All rights reserved.

00052618 19990601152B0267 (USE FORMAT 7 FOR FULLTEXT)

SmartASIC Introduces TFT LCD Display Controller for 16-19" Screens; Lowers Cost of Computer Monitors
and Projectors

Tuesday , June 1, 1999 12:01 EDT

Word Count: 735

Company Names: STANFORD RESOURCES INC

Geographic Names: CALIFORNIA; AMERICAS; NORTH AMERICA; USA

Product Names: COMPUTERS; OPTOELECTRONICS; ELECTRONICS INDUSTRY

Event Names: TECHNOLOGY DEVELOPMENT

4/8/27 (Item 5 from file: 610)

Business Wire

(c) 2008 Business Wire. All rights reserved.

00021764 1999087B0003 (USE FORMAT 7 FOR FULLTEXT)

Cypress Nears Close of IC WORKS Acquisition; Shareholder, Regulatory Approval Process Now Complete

Sunday , March 28, 1999 21:20 EST

Word Count: 949

Company Names: SEMICONDUCTOR HOLDING BV; CYPRESS SEMICONDUCTOR CORP; IC WORKS INC; ICW GROUP PLC; ICW INC; FEDERAL TRADE COMMISSION; NATIONAL SEMICONDUCTOR CORP; SAMSUNG; COMPUTER PRODUCTS INC

Geographic Names: CALIFORNIA; NEW YORK; USA; AMERICAS; NORTH AMERICA

Product Names: MERGERS AND ACQUISITIONS; MICROCHIPS; REGULATION; SEMICONDUCTORS; STOCKS AND SHARES; CORPORATE; ELECTRONIC COMPONENTS; ELECTRONICS INDUSTRY; INSTITUTIONS; FINANCIAL SERVICES; INVESTMENT

Event Names: INVESTMENT; MANUFACTURING AND PRODUCTION; MERGERS AND ACQUISITIONS; REGULATION; STOCKS AND SHARES

4/8/28 (Item 1 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0949984 BW0132

SUN MICROSYSTEMS : Sun Releases Beta Java 2 Platform Optimized for Solaris to Help ISVs Run Enterprise-Class Java Applications

December 09, 1998

Byline: Business Editors & Computer Writers

Word Count: 919

4/8/29 (Item 2 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0929128 BW0119

SUN MICROSYSTEMS 6 : Built for Business - JDK 1.2 for the Solaris 7 Operating Environment; Sun's Java Applications for the World's Strongest Operating Environment - Three Times Faster Than NT

October 27, 1998

Byline: Business Editors/High Tech Writers
Word Count: 855

4/8/30 (Item 3 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0901797 BW0360

TERA COMPUTING : Tera Computer Company Unveils Supercomputer Roadmap Providing a Future for SGI/Cray T90 Users

September 01, 1998

Byline: Business Editors/Computer Writers
Word Count: 1655

4/8/31 (Item 4 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0901792 BW0358

CQN TERA COMPUTING : Tera Computer Corrects and Replaces Previous Product Announcement, BW285, TERA-COMPUTER

September 01, 1998

Byline: Business Editors/Computer Writers
Word Count: 1681

4/8/32 (Item 5 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0885263 BW1093

ROCKWELL SEMICONDUCTOR : Rockwell Semiconductor Systems is First To Take a Single T1/E1/J1 Frammer Chip to Octal Density

July 27, 1998

Byline: Business Editors and High-Tech Writers

Word Count: 1293

4/8/33 (Item 6 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0873323 BW1078

PERVASIVE : Pervasive and Synchrologic Team To Create Shrink-Wrap Mobile Database Solution;
Relationship Adds Leading Synchronization Technology to Pervasive's Ultralight Database

June 29, 1998

Byline: Business/Technology Editors

Word Count: 585

4/8/34 (Item 7 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0830013 BW1130

COMPAQ 2 : Compaq Introduces Ultimate Video Conferencing Kit and High-Capacity Diskette Drive for
Portable PCs

April 02, 1998

Byline: Business/Technology Editors

Word Count: 1204

4/8/35 (Item 8 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0830009 BW1127

COMPAQ : Compaq Unveils Powerful Armada 7800 Notebook PC Featuring Intel's Mobile Pentium II
Processor

April 02, 1998

Byline: Business/Technology Editors

Word Count: 2984

4/8/36 (Item 9 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0774312 BW1463

HITACHI HOME ELEC : Hitachi Announces Next Generation Handheld PC With Microsoft Windows CE 2.0

November 17, 1997

Byline: Business Editors & Technology Writers

Word Count: 1172

4/8/37 (Item 10 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0763830 BW1178

SMART MODULAR TECH : SMART Modular Technologies Announces High-Density Registered SDRAM Modules for High-end Systems

October 27, 1997

Byline: Business Editors/Computer Writers

Word Count: 602

4/8/38 (Item 11 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0668277 BW1376

VTEL : VTEL Assigned Patent for Multipoint Videoconference Technology

February 03, 1997

Byline: Business Editors

Word Count: 342

4/8/39 (Item 12 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0496657 BW1028

NEC ELECTRONICS : NEC Electronics Inc. Debuts High-Density ASIC With CBA Architecture; CMOS-8LHD Ideal for High-Integration Designs

June 26, 1995

Byline: Business Editors

Word Count: 956

4/8/40 (Item 13 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0333844 BW625

MOTOROLA : Motorola unveils next-generation 8-bit microcontroller architecture

May 10, 1993

Byline: Business Editors and Computers Writers

Word Count: 1431

4/8/41 (Item 1 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

02384458 Supplier Number: 60805620 (Use Format 7 Or 9 For FULL TEXT)

Legato Updates Three Products for Windows 2000.(Product Announcement)

March 8 , 2000

Word Count: 231 Line Count: 00022

Company Names: Legato Systems Inc.--Product introduction

Geographic Codes/Names: 1USA United States

Descriptors: Backup software; Network software; Networking software product introduction

Event Codes/Names: 336 Product introduction

Product/Industry Names: 7372620 (Network Software)

SIC Codes: 7372 Prepackaged software

NAICS Codes: 51121 Software Publishers

Ticker Symbols: lgto

Trade Names: Legato NetWorker 5.7 (Backup software)--Product introduction; Legato Octopus 4.0 (Backup software)--Product introduction; Legato Cluster Enterprise 4.5.1 (Network software)--Product introduction

File Segment: CD File 275

4/8/42 (Item 2 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

02377084 Supplier Number: 59624577 (Use Format 7 Or 9 For FULL TEXT)
new products.

Feb , 2000

Word Count: 4177 Line Count: 00357

File Segment: CD File 275

4/8/43 (Item 3 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

02376837 Supplier Number: 59664505 (Use Format 7 Or 9 For FULL TEXT)

Cell phones strike back : Choices between handheld PCs and cellular phones get tougher as wireless device development accelerates.(Industry Trend or Event)

Feb 28 , 2000

Word Count: 571 Line Count: 00050

Company Names: Phone.com Inc.--Product development; Intel Corp.--Product development

Geographic Codes/Names: 1USA United States

Descriptors: Industry trend; Smart phone; Flash memory; Internet

Event Codes/Names: 336 Product introduction

Product/Industry Names: 7372681 (Internet Access Software); 4811500 (Specialized Telecommunication Services); 3662166 (Cellular Telephones); 3573221 (Computer RAM)

NAICS Codes: 51121 Software Publishers; 51331 Wired Telecommunications Carriers; 33422 Radio and Television Broadcasting and Wireless Communications Equipment Manufacturing; 334413 Semiconductor and Related Device Manufacturing

File Segment: CD File 275

4/8/44 (Item 4 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

02365011 Supplier Number: 58925511 (Use Format 7 Or 9 For FULL TEXT)

Motorola Reports Higher Fourth-Quarter, Full-Year Sales and Earnings.(Company Financial Information)

Jan 24 , 2000

Word Count: 3240 Line Count: 00447

Company Names: Motorola Inc.--Finance

Geographic Codes/Names: 1USA United States

Descriptors: Electronics industry; Company sales/revenue; Company earnings/profit

Event Codes/Names: 830 Sales, profits & dividends

Product/Industry Names: 3601000 (Electronics)

NAICS Codes: 3359 Other Electrical Equipment and Component Manufacturing

File Segment: CD File 275

4/8/45 (Item 5 from file: 275)

Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02277748 Supplier Number: 54090766 (Use Format 7 Or 9 For FULL TEXT)
HOOKED ON PLACEBOS.

April , 1999
Word Count: 1623 Line Count: 00130
File Segment: CD File 275

4/8/46 (Item 6 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02268125 Supplier Number: 53741443 (Use Format 7 Or 9 For FULL TEXT)
A Wireless Entanglement.(Ericsson's Bluetooth specification for wireless 'personal area networks')(Company Business and Marketing)(Column)

March 9 , 1999
Word Count: 829 Line Count: 00067
Company Names: LM Ericsson Telefon AB--Standards
Geographic Codes/Names: 1USA United States
Descriptors: Wireless network; Standard; Company technology development
Event Codes/Names: 350 Product standards, safety, & recalls
Product/Industry Names: 3662100 (Communications Equipment ex Broadcast)
SIC Codes: 3660 Communications Equipment
NAICS Codes: 33429 Other Communications Equipment Manufacturing
File Segment: CD File 275

4/8/47 (Item 7 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02206808 Supplier Number: 21004107 (Use Format 7 Or 9 For FULL TEXT)
Chips: Rockwell Semiconductor Systems is First To Take a Single T1/E1/J1 Frammer Chip to Octal Density.(the RS8398 chip provides a single universal octal solution for physical layer termination of multiplexed voice and data traffic)(Product Announcement)

August 3 , 1998
Word Count: 1248 Line Count: 00107
Company Names: Rockwell Semiconductor Systems--Product introduction
Descriptors: Multiplexer; Networking Hardware Product Introduction
Product/Industry Names: 3674182 (Multiplexer Circuits)
SIC Codes: 3674 Semiconductors and related devices
Trade Names: Rockwell Semiconductor RS8398 (Multiplexer)--Product introduction
File Segment: CD File 275

4/8/48 (Item 8 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02192403 Supplier Number: 19758643 (Use Format 7 Or 9 For FULL TEXT)
Sharing memory with memory-mapped files. (Technology Tutorial)

Oct , 1997
Word Count: 2723 Line Count: 00230

Special Features: program; illustration
Descriptors: Programming Tutorial; UNIX
File Segment: CD File 275

4/8/49 (Item 9 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02159872 Supplier Number: 20480777 (Use Format 7 Or 9 For FULL TEXT)
New Notebooks: Compaq Unveils Powerful Armada 7800 Notebook PC Featuring Intel's Mobile Pentium II Processor.(Product Announcement)

April 6 , 1998
Word Count: 2953 Line Count: 00248
Descriptors: Hardware Product Introduction; Pentium II-Based Notebook
Product/Industry Names: 3573141 (Intel-Compatible Notebook Computers)
SIC Codes: 3571 Electronic computers
Trade Names: Compaq Armada 7800 (Pentium II-based notebook)--Product introduction
File Segment: CD File 275

4/8/50 (Item 10 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02123441 Supplier Number: 20027758 (Use Format 7 Or 9 For FULL TEXT)
Windows CE: Hitachi announces next generation Handheld PC with Microsoft Windows CE 2.0. (Hitachi Handheld PC)(Product Announcement)

Nov 24 , 1997
Word Count: 983 Line Count: 00087
Company Names: Hitachi Home Electronics (America) Inc.--Product introduction
Descriptors: Hardware Product Introduction; Personal Digital Assistant
Product/Industry Names: 3573160 (Personal Digital Assistants)
SIC Codes: 3571 Electronic computers
Trade Names: Hitachi Handheld PC (Personal digital assistant)--Product introduction
File Segment: CD File 275

4/8/51 (Item 11 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02075514 Supplier Number: 19528096 (Use Format 7 Or 9 For FULL TEXT)
Memory overload: Making sense of RAM. (Technology Information)

June 23 , 1997
Word Count: 557 Line Count: 00046

Special Features: illustration; table
Company Names: Apple Computer Inc.--Products
Descriptors: Technology Overview; RAM; DRAM; SRAM; Microcomputer Industry
Product/Industry Names: 3674125 (Random Access Memory Circuits)
SIC Codes: 3674 Semiconductors and related devices
Ticker Symbols: AAPL
File Segment: CD File 275

4/8/52 (Item 12 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02013777 Supplier Number: 18875632 (Use Format 7 Or 9 For FULL TEXT)
Chips: C-Cube delivers real-time digital video encoding to consumer PC applications with debut of low-cost encoder chip family; C-Cube's CLM41xx. (C-Cube Microsystems CLM4100 MPEG-1 encoders)(Product Announcement)

Oct 28 , 1996
Word Count: 1181 Line Count: 00097
Company Names: C-Cube Microsystems Inc.--Product introduction
Descriptors: Hardware Product Introduction; Video Processing Equipment
Product/Industry Names: 3573250 (Computer Optical & Graphics Equip); 3662650 (Image Processing Equip)
SIC Codes: 3577 Computer peripheral equipment, not elsewhere classified
Ticker Symbols: CUBE
Trade Names: C-Cube Microsystems CLM4100 (Video processing equipment)--Product introduction
File Segment: CD File 275

4/8/53 (Item 13 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
01966984 Supplier Number: 18564817
Reporting against large databases. (the role of server-based reporting engines) (Technology Information)

August , 1996
Word Count: 2211 Line Count: 00183

Special Features: illustration; chart

Descriptors: Technology Overview; DBMS; Report Generation Software; Database Design; Data Warehousing; Client/Server Architecture

SIC Codes: 7372 Prepackaged software

File Segment: CD File 275

4/8/54 (Item 14 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01944414 Supplier Number: 18371553 (Use Format 7 Or 9 For FULL TEXT)

Pentium Classic: still the one. (Overview of evaluations 101 Pentium-based systems) (individual evaluation articles searchable under "Pentium Classic: Still the One)(includes related articles on the editors' choices, Pentium vs. Pentium Pro performance, reading the Service & Reliability boxes, Pentium PC features, benchmark test results, purchasing guidelines, price/performance index, and summary of features) (Hardware Review)(Evaluation)(Cover Story)

June 25 , 1996

Word Count: 7989 Line Count: 00596

Special Features: illustration; photograph; table; chart; graph

Company Names: Dell Computer Corp.--Products; Micron Electronics Inc.--Products

Descriptors: Hardware Multiproduct Review; Pentium-Based System

SIC Codes: 3571 Electronic computers

Ticker Symbols: DELL

Trade Names: Dell Computer Dell Dimension XPS P133c (Pentium-based system)-- Evaluation; Micron Electronics P133 Millennia (Pentium-based system)-- Evaluation; Micron Electronics P166 Millennia (Pentium-based system)-- Evaluation

File Segment: CD File 275

4/8/55 (Item 15 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01938349

Supplier Number: 18295599 (Use Format 7 Or 9 For FULL TEXT)

New RAM burns rubber. (synchronous DRAM) (Technology Information)

May 10 , 1996

Word Count: 430 Line Count: 00035

Company Names: Dell Computer Corp.--Products

Descriptors: DRAM; Microcomputer Industry; Technology Overview

SIC Codes: 3571 Electronic computers

Ticker Symbols: DELL

File Segment: CD File 275

4/8/56 (Item 16 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01902071 Supplier Number: 17946189 (Use Format 7 Or 9 For FULL TEXT)

A new memory system design for commercial and technical computing products. (HP's J/K-class memory system design) (Technology Information)

Feb , 1996

Word Count: 5557 Line Count: 00429

Special Features: illustration; chart

Descriptors: Technology Overview; System Design; Semiconductor Memory

SIC Codes: 3674 Semiconductors and related devices

File Segment: CD File 275

4/8/57 (Item 17 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01822435 Supplier Number: 17115295 (Use Format 7 Or 9 For FULL TEXT)

Philips cooking up full menu. (Philips Semiconductors reorganizing, forming new product, manufacturing and acquisition plans)

June 26 , 1995

Word Count: 2181 Line Count: 00181

Special Features: illustration; photograph; chart

Company Names: Philips Semiconductors--Planning

Descriptors: Company Operations; Company Restructuring/Company Reorganization; Company Business And Marketing

File Segment: CD File 275

4/8/58 (Item 18 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01819197 Supplier Number: 17365620 (Use Format 7 Or 9 For FULL TEXT)

When CPUs share.(Symmetric Multiprocessing and its hidden problems)

June 16 , 1995

Word Count: 1291 Line Count: 00105

Descriptors: CPU; Processor Architecture; Multiprocessing; Technology Information ; Technology Development

File Segment: CD File 275

4/8/59 (Item 19 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
01777652 Supplier Number: 16864614 (Use Format 7 Or 9 For FULL TEXT)
Bottlestopper: deepspace Windows. (The Soft Side)(Column)

April , 1995
Word Count: 795 Line Count: 00061
File Segment: CD File 275

4/8/60 (Item 20 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
01687496 Supplier Number: 16056226 (Use Format 7 Or 9 For FULL TEXT)
Solid for midrange LANs; HP's NetServer LF fits the bill for medium-sized workgroups, with a few caveats. (HP's network server) (includes related article on testing methodology) (PC Week LABS: First Look) (PC Week Netweek) (Hardware Review) (Evaluation)

June 20 , 1994
Word Count: 1613 Line Count: 00127

Special Features: illustration; table; graph
Company Names: Hewlett-Packard Co.--Products
Descriptors: Evaluation; File Server
Product/Industry Names: 3573115 (Microcomputers)
SIC Codes: 3571 Electronic computers
Ticker Symbols: HWP
Trade Names: HP NetServer LF (Intel-compatible system)--Evaluation
File Segment: CD File 275

4/8/61 (Item 21 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
01679613 Supplier Number: 15313342 (Use Format 7 Or 9 For FULL TEXT)
Advances in parallel computing for reactor analysis and safety. (research underway at Argonne National Laboratory and other facilities for nuclear power plant simulation, improvements under parallel architecture) (High Performance Computing)

April , 1994
Word Count: 6403 Line Count: 00537

Special Features: illustration; graph; chart
Descriptors: Simulation; Theoretical Research; Argonne National Laboratory; Nuclear energy; Control Systems; Parallel processing

SIC Codes: 3443 Fabricated plate work (boiler shops); 4911 Electric services
File Segment: CD File 275

4/8/62 (Item 22 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01667427 Supplier Number: 15066256 (Use Format 7 Or 9 For FULL TEXT)

At your service. (AST Research's Manhattan SMP P/60, Compaq's ProLiant 2000 5/66 M4200A, HP's NetServer 5/60 LM, Unisys' PW2 Advantage Plus 5608, Wyse Technology's 7000i 760MP application servers) (includes five related articles on recent major developments, top-rated Compaq ProLiant 2000 system, suitability to task, price/performance ratio and benchmark testing) (Hardware Review) (Evaluation)

March 15 , 1994

Word Count: 9374 Line Count: 00733

Special Features: illustration; photograph; table; graph

Company Names: AST Research Inc.--Products; Compaq Computer Corp.--Products; Hewlett-Packard Co.--Products; Unisys Corp.--Products; Wyse Technology Inc. --Products

Descriptors: Evaluation; File Server

SIC Codes: 3571 Electronic computers; 3577 Computer peripheral equipment, not elsewhere classified; 3575 Computer terminals

Ticker Symbols: ASTA; UIS; HWP; CPQ; WYS

Trade Names: AST Research Manhattan SMP P/60 (Pentium-based system)--evaluation; Compaq ProLiant 2000

5/66 M4200A (Pentium-based system)--evaluation; HP NetServer 5/60 LM (Pentium-based system)--evaluation;

Unisys PW2 Advantage Plus 5608 (Pentium-based system)--evaluation; Wyse Technology 7000i 760MP (486-based system)--evaluation

File Segment: CD File 275

4/8/63 (Item 23 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01623571 Supplier Number: 14468926 (Use Format 7 Or 9 For FULL TEXT)

NetWare 4.0 for database developers. (Novell's network operating system) (Software Review) (includes related articles on changes in version 4.0 and problems that arose when installing maintenance version 4.01) (Evaluation)

Oct , 1993

Word Count: 5608 Line Count: 00432

Special Features: illustration; table

Company Names: Novell Inc.--Products

Descriptors: Evaluation; Network Operating System

SIC Codes: 7372 Prepackaged software

Ticker Symbols: NOVL

Trade Names: NetWare 4.0 (Network operating system)--evaluation

File Segment: CD File 275

4/8/64 (Item 24 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01614381 Supplier Number: 14192541 (Use Format 7 Or 9 For FULL TEXT)

Bridging the gap between structured analysis and structured design for real-time systems. (includes related article on principles of structured analysis and design) (Technical)

August , 1993

Word Count: 4551 Line Count: 00384

Special Features: illustration; chart; table

Descriptors: Real-Time System; Structured Design Techniques; Systems Analysis; System Design; New Technique; Image Processing; Medical Diagnosis

SIC Codes: 7372 Prepackaged software

File Segment: CD File 275

4/8/65 (Item 25 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01586747 Supplier Number: 13393030 (Use Format 7 Or 9 For FULL TEXT)

Take charge of your network! (tips for managing local area networks)(includes related articles on fine-tuning DOS, keeping NetWare drivers up to date, securing print servers, invoking password protection on screen saver, determining whether user is logged on to network, managing print jobs) (Tutorial)

March , 1993

Word Count: 6005 Line Count: 00465

Descriptors: LAN; Network Management; Tutorial; Management of EDP; Data security

Operating Platform: NetWare

File Segment: CD File 275

4/8/66 (Item 26 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01581233 Supplier Number: 13085423 (Use Format 7 Or 9 For FULL TEXT)

Utilities put the power in PBs. (PowerBook utility collections) (Product Watch)

Jan 4 , 1993

Word Count: 821 Line Count: 00066

Company Names: Apple Computer Inc.--Products; Connectix Corp.--Products; Symantec Corp.--Products

Descriptors: Desktop Utility; Software Design; Software Selection; Computer software industry

SIC Codes: 7372 Prepackaged software; 3571 Electronic computers

Ticker Symbols: AAPL; SYMC

Trade Names: Apple Macintosh PowerBook (Notebook computer)--Computer programs; Connectix PowerBook Utilities (Operating system enhancement)-- Design and construction; Norton Essentials for PowerBook (Operating system enhancement)--Design and construction
Operating Platform: Apple Macintosh
File Segment: CD File 275

4/8/67 (Item 27 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
01553337 Supplier Number: 13377724 (Use Format 7 Or 9 For FULL TEXT)
DSP chip set aimed at real time. (Sharp Microelectronics' LH9124 processor and LH9320 address generator)
(Product Announcement)

Oct , 1992
Word Count: 557 Line Count: 00045

Special Features: illustration; table
Company Names: Sharp Microelectronics Technology Inc.--Product introduction
Descriptors: Product Introduction; Digital Signal Processor; Integrated Circuits; Chip Set; Real-Time System
SIC Codes: 8731 Commercial physical research; 3674 Semiconductors and related devices
Trade Names: Sharp LH9124 (Digital signal processor)--Product introduction; Sharp LH9320 (Semiconductor device)--Product introduction
File Segment: CD File 275

4/8/68 (Item 28 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
01530956 Supplier Number: 12515725 (Use Format 7 Or 9 For FULL TEXT)
Scalability. (Ultracomputers: a Terflop Before its Time)

August , 1992
Word Count: 5648 Line Count: 00467

Special Features: illustration; chart
Descriptors: Scales; Performance Improvement; Optimization; Expandability; Processor Speed; Size; Computers; Generations of Computers; Parallel Processing; Computer industry
SIC Codes: 3571 Electronic computers
File Segment: CD File 275

4/8/69 (Item 29 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
01530423 Supplier Number: 12528179 (Use Format 7 Or 9 For FULL TEXT)
Tony Agnello on: DSPs in multiprocessing. (digital signal processors) (Technology Viewpoint) (Column)

August , 1992

Word Count: 2454 Line Count: 00207

Descriptors: Multiprocessing; Digital Signal Processor; Technology; Integrated Circuits; Integrated Circuit Cards; Circuit Design

SIC Codes: 3674 Semiconductors and related devices

File Segment: CD File 275

4/8/70 (Item 30 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01528138 Supplier Number: 12509263 (Use Format 7 Or 9 For FULL TEXT)

Vendors reply to frame relay 20 questions. (Communications Management) (Column)

July , 1992

Word Count: 996 Line Count: 00086

Descriptors: Frame Relay; Telecommunications Services Industry; Communications Management; Purchases; Hardware Selection; Communications Equipment; Packet Switch; LAN; Data communications

SIC Codes: 4800 COMMUNICATION; 3661 Telephone and telegraph apparatus; 3660 Communications Equipment

Operating Platform: Frame Relay

File Segment: CD File 275

4/8/71 (Item 31 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01512579 Supplier Number: 12230349 (Use Format 7 Or 9 For FULL TEXT)

VMS POSIX's challenge to UNIX. (includes related article on X/Open standards)

May , 1992

Word Count: 3161 Line Count: 00259

Special Features: illustration; chart; table

Descriptors: POSIX Standard; Operating System; Software Design; Standard; Systems Software; UNIX; Competition; Trends

File Segment: CD File 275

4/8/72 (Item 32 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01502698 Supplier Number: 11957965 (Use Format 7 Or 9 For FULL TEXT)

Understanding video displays: from CRTs to shadow masks. (cathode ray tubes) (Tech Section; includes related article on video troubleshooting) (Tutorial)

March , 1992

Word Count: 2830 Line Count: 00230

Special Features: illustration; table

Descriptors: CRT Display; Color; Tutorial; Monitors

File Segment: CD File 275

4/8/73 (Item 33 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01438519 Supplier Number: 10936554 (Use Format 7 Or 9 For FULL TEXT)

What line scan vision systems can do for you. (includes related article on line scan vision equipment)

May , 1991

Word Count: 2348 Line Count: 00186

Special Features: illustration; photograph; chart

Descriptors: Vision; Scanner; Image Processing; Manufacturing; Measurement; Inspection; Microcomputer;

Software; Cameras; Line Monitors

File Segment: CD File 275

4/8/74 (Item 34 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01415150 Supplier Number: 09240412 (Use Format 7 Or 9 For FULL TEXT)

Microsoft touts multimedia PC; Windows extensions to developers.

Jan 1 , 1991

Word Count: 1835 Line Count: 00149

Company Names: Microsoft Corp.--Product introduction

Descriptors: Multimedia Technology; Application Development Software; Product Development; Applications

Programming; GUI; Market Analysis

SIC Codes: 7372 Prepackaged software

Ticker Symbols: MSFT

Operating Platform: MS Windows

File Segment: CD File 275

4/8/75 (Item 35 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01383293 Supplier Number: 09483055 (Use Format 7 Or 9 For FULL TEXT)

Software products support process and manufacturing industries. (product announcement)

Sept , 1990

Word Count: 8026 Line Count: 00651

Special Features: illustration; table

Company Names: Digital Equipment Corp.--Product introduction

Descriptors: Connectivity; Computer-Integrated Manufacturing; Applications; Manufacturing; Process Control;

Product Introduction; Software Packages

SIC Codes: 7373 Computer integrated systems design

Ticker Symbols: DEC

Trade Names: Process/Lab Integration Set (CIM software)--Product introduction; DEComni/VMS (CIM software)-
-Product introduction

File Segment: CD File 275

4/8/76 (Item 36 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01355132 Supplier Number: 08330674 (Use Format 7 Or 9 For FULL TEXT)

Retrospective on DACNOS. (prototype Distributed Academic Computing Network Operating System)

April , 1990

Word Count: 9077 Line Count: 00744

Descriptors: Multivendor Systems; Distributed Processing; Network Operating System; Prototype; LAN;
Education

File Segment: AI File 88

4/8/77 (Item 37 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01310577 Supplier Number: 07745360 (Use Format 7 Or 9 For FULL TEXT)

Technical correspondence.

Oct , 1989

Word Count: 15663 Line Count: 01233

File Segment: CD File 275

4/8/78 (Item 38 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01286500 Supplier Number: 07298051 (Use Format 7 Or 9 For FULL TEXT)

Booting up and shutting down your system. (Daemons and Dragons) (column)

Jan , 1989

Word Count: 2179 Line Count: 00172

Descriptors: UNIX; Tutorial; Booting; Systems Analysis; Troubleshooting; Time Sharing

SIC Codes: 7372 Prepackaged software

Operating Platform: Unix

File Segment: CD File 275

4/8/79 (Item 39 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01249600 Supplier Number: 06746933 (Use Format 7 Or 9 For FULL TEXT)

Support for tightly coupled processors.

May , 1988

Word Count: 3216 Line Count: 00261

Company Names: Stellar Computer Inc.--Products

Descriptors: UNIX; Coupling; Microprocessor; Integrated Circuits

SIC Codes: 3674 Semiconductors and related devices

Trade Names: Stellar Computer GS1000 (Workstation)

File Segment: CD File 275

4/8/80 (Item 40 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01204969 Supplier Number: 06329711 (Use Format 7 Or 9 For FULL TEXT)

For cost-performance, partition RISC system on bus parameters.

Nov 12 , 1987

Word Count: 3158 Line Count: 00255

Special Features: illustration; chart

Company Names: VLSI Technology Inc.--Innovations

SIC Codes: 3571 Electronic computers

File Segment: TI File 148

4/8/81 (Item 41 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01204424 Supplier Number: 04729605 (Use Format 7 Or 9 For FULL TEXT)

Data-acquisition system fits on a smart peripheral chip.

March 5 , 1987

Word Count: 2936 Line Count: 00226

Special Features: illustration; photograph

SIC Codes: 3674 Semiconductors and related devices
File Segment: TI File 148

4/8/82 (Item 1 from file: 621)
Gale Group New Prod.Annou.(R)
(c) 2008 Gale/Cengage. All rights reserved.
02282849 Supplier Number: 58611094 (USE FORMAT 7 FOR FULLTEXT)
Motorola Reports Higher Fourth-Quarter, Full-Year Sales and Earnings.

Jan 17 , 2000
Word Count: 4190
Publisher Name: Business Wire
Company Names: *Iridium L.L.C.; Motorola Inc.
Geographic Names: *1USA (United States)
Product Names: *3601000 (Electronics); 3662130 (Satellite Communications Systems)
Industry Names: BUS (Business, General); BUSN (Any type of business)
SIC Codes: 3663 (Radio & TV communications equipment); 3670 (Electronic Components and Accessories)
NAICS Codes: 3359 (Other Electrical Equipment and Component Manufacturing); 33422 (Radio and Television Broadcasting and Wireless Communications Equipment Manufacturing)
Ticker Symbols: MOT

4/8/83 (Item 2 from file: 621)
Gale Group New Prod.Annou.(R)
(c) 2008 Gale/Cengage. All rights reserved.
02118226 Supplier Number: 55148447 (USE FORMAT 7 FOR FULLTEXT)
Antex Electronics Introduces New Model BX-44 to Broadcaster Series of Digital Audio Cards.

July 14 , 1999
Word Count: 687
Publisher Name: Business Wire
Geographic Names: *1USA (United States)
Product Names: *3573293 (Computer Graphics, Sound and Video Processors)
Industry Names: BUS (Business, General); BUSN (Any type of business)
SIC Codes: 3577 (Computer peripheral equipment, not elsewhere classified)
NAICS Codes: 334119 (Other Computer Peripheral Equipment Manufacturing)

4/8/84 (Item 3 from file: 621)
Gale Group New Prod.Annou.(R)
(c) 2008 Gale/Cengage. All rights reserved.
01562639 Supplier Number: 47921166 (USE FORMAT 7 FOR FULLTEXT)
PNY's SDRAM Memory Upgrades Support the Latest IBM Aptiva Multimedia Deskside Systems

August 19 , 1997
Word Count: 542

Publisher Name: PR Newswire Association, Inc.
Company Names: *PNY Technologies Inc.
Event Names: *336 (Product introduction)
Geographic Names: *1USA (United States)
Product Names: *3674126 (IC Memory Chips)
Industry Names: BUS (Business, General); BUSN (Any type of business)
NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)

4/8/85 (Item 4 from file: 621)
Gale Group New Prod.Annou.(R)
(c) 2008 Gale/Cengage. All rights reserved.
01441361 Supplier Number: 46813394 (USE FORMAT 7 FOR FULLTEXT)
C-Cube Delivers Real-Time Digital Video Encoding to Consumer PC Applications with Introduction of Low-Cost Encoder Chip Family; C-Cube's CLM41xx Product Family Transforms Digital Video on the PC Into an Active Data Type for Internet, Desktop Video, and CD-Authoring Applications.

Oct 21 , 1996
Word Count: 1111
Publisher Name: Business Wire
Company Names: *C-Cube Microsystems Inc.
Event Names: *330 (Product information)
Geographic Names: *1USA (United States)
Product Names: *3662600 (Signal Processing Equipment)
Industry Names: BUS (Business, General); BUSN (Any type of business)
NAICS Codes: 33429 (Other Communications Equipment Manufacturing)
Ticker Symbols: CUBE

4/8/86 (Item 5 from file: 621)
Gale Group New Prod.Annou.(R)
(c) 2008 Gale/Cengage. All rights reserved.
01138217 Supplier Number: 41212126 (USE FORMAT 7 FOR FULLTEXT)
SWITCHMODE CONTROLLER RUNS FROM 12V BATTERY

March 6 , 1990
Word Count: 442
Publisher Name: Various
Company Names: *Teledyne Components Inc.
Event Names: *330 (Product information)
Geographic Names: *1USA (United States); 1U9CA (California)

Product Names: *3674156 (IC Voltage Multipliers & Regulators)
Industry Names: BUS (Business, General); BUSN (Any type of business)
NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)
Trade Names: TSC9112

4/8/87 (Item 6 from file: 621)

Gale Group New Prod.Annou.(R)

(c) 2008 Gale/Cengage. All rights reserved.

01087128 Supplier Number: 40532414 (USE FORMAT 7 FOR FULLTEXT)

TP3420/ST5420A S INTERFACE DEVICE

Oct 4 , 1988

Word Count: 465

Publisher Name: Various

Company Names: *National Semiconductor Corp.; SGS Thomson Microelectronics S.R.L.

Event Names: *380 (Strategic alliances)

Geographic Names: *1USA (United States); 1U2NY (New York)

Product Names: *3674199 (ICs by Function NEC)

Industry Names: BUS (Business, General); BUSN (Any type of business)

NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)

Ticker Symbols: NSM

Trade Names: TP3420/ST5420S

4/8/88 (Item 7 from file: 621)

Gale Group New Prod.Annou.(R)

(c) 2008 Gale/Cengage. All rights reserved.

01087099 Supplier Number: 40532385 (USE FORMAT 7 FOR FULLTEXT)

NATIONAL SEMICONDUCTOR AND SGS-THOMSON MICROELECTRONICS ANNOUNCE THEIR
FIRST JOINTLY DEVELOPED ISDN PRODUCTS

Oct 4 , 1988

Word Count: 905

Publisher Name: Various

Company Names: *National Semiconductor Corp.; SGS Thomson Microelectronics S.R.L.

Event Names: *380 (Strategic alliances)

Geographic Names: *1USA (United States); 1U2NY (New York)

Product Names: *3674199 (ICs by Function NEC)

Industry Names: BUS (Business, General); BUSN (Any type of business)

NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)

Ticker Symbols: NSM

Trade Names: TP3420/ST5420; TP3076/ST5076

4/8/89 (Item 1 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

04540689 Supplier Number: 58946674 (USE FORMAT 7 FOR FULLTEXT)

Fujitsu develops high performance graphics display controller.

Jan 24 , 2000

Word Count: 1059

Publisher Name: M2 Communications Ltd.

Company Names: *Fujitsu Laboratories Ltd.

Geographic Names: *9JAPA (Japan)

Product Names: *3573293 (Computer Graphics, Sound and Video Processors); 3674000 (Semiconductor Devices)

Industry Names: BUSN (Any type of business); INTL (Business, International)

SIC Codes: 3577 (Computer peripheral equipment, not elsewhere classified); 3674 (Semiconductors and related devices)

NAICS Codes: 334119 (Other Computer Peripheral Equipment Manufacturing); 334413 (Semiconductor and Related Device Manufacturing)

4/8/90 (Item 2 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

04438582 Supplier Number: 55814615 (USE FORMAT 7 FOR FULLTEXT)

NOTEBOOK.

Sept 20 , 1999

Word Count: 2428

Publisher Name: Warren Publishing, Inc.

Company Names: *Matsushita Electric Industrial Company Ltd.

Event Names: *443 (New capacity, new plant construction)

Geographic Names: *9JAPA (Japan)

Product Names: *3679582 (Liquid Crystal Displays); 3600000 (Electrical & Electronic Equip)

Industry Names: BUSN (Any type of business); ELEC (Electronics)

NAICS Codes: 334419 (Other Electronic Component Manufacturing); 335 (Electrical Equipment, Appliance, and Component Manufacturing)

4/8/91 (Item 3 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

04040046 Supplier Number: 53398843 (USE FORMAT 7 FOR FULLTEXT)

SUN MICROSYSTEMS: Sun releases Beta Java 2 platform optimized for Solaris.

Dec 10 , 1998

Word Count: 926

Publisher Name: M2 Communications

Industry Names: BUSN (Any type of business); INTL (Business, International)

4/8/92 (Item 4 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

04000998 Supplier Number: 53145247 (USE FORMAT 7 FOR FULLTEXT)

-SUN MICROSYSTEMS: Built for business -- JDK 1.2 for the Solaris 7 operating environment.

Oct 28 , 1998

Word Count: 864

Publisher Name: M2 Communications

Company Names: *Sun Microsystems Inc.

Geographic Names: *1USA (United States)

Product Names: *3573000 (Computers & Peripherals); 7372513 (Application Development Software)

Industry Names: BUSN (Any type of business); INTL (Business, International)

NAICS Codes: 334111 (Electronic Computer Manufacturing); 51121 (Software Publishers)

4/8/93 (Item 5 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

03935574 Supplier Number: 50213530 (USE FORMAT 7 FOR FULLTEXT)

Chips: Rockwell Semiconductor Systems is First To Take a Single T1/E1/J1 Frammer Chip to Octal Density

August 3 , 1998

Word Count: 1163

Publisher Name: EDGE Publishing

Company Names: *Rockwell Semiconductor Systems

Event Names: *336 (Product introduction)

Geographic Names: *1USA (United States)

Product Names: *3674124 (Microprocessor Chips)

Industry Names: BUSN (Any type of business); TELC (Telecommunications)

NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)

4/8/94 (Item 6 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

03861927 Supplier Number: 48408402 (USE FORMAT 7 FOR FULLTEXT)

New Notebooks: Compaq Unveils Powerful Armada 7800 Notebook PC Featuring Intel's Mobile Pentium II Processor

April 6 , 1998

Word Count: 2770

Publisher Name: EDGE Publishing

Company Names: *Compaq Computer Corp.

Event Names: *330 (Product information)

Geographic Names: *1USA (United States)

Product Names: *3573140 (Notebook Computers)

Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation); TELC (Telecommunications)

NAICS Codes: 334111 (Electronic Computer Manufacturing)

Ticker Symbols: CPQ

4/8/95 (Item 7 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

03859812 Supplier Number: 48401089 (USE FORMAT 7 FOR FULLTEXT)

-COMPAQ: Compaq unveils powerful Armada 7800 Notebook PC featuring Intel's Mobile Pentium II Processor

April 3 , 1998

Word Count: 2286

Publisher Name: M2 Communications

Industry Names: BUSN (Any type of business); INTL (Business, International)

4/8/96 (Item 8 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

03859811 Supplier Number: 48401088 (USE FORMAT 7 FOR FULLTEXT)

-COMPAQ: Compaq introduces ultimate videoconferencing kit & high-capacity diskette drive for portable PCs

April 3 , 1998

Word Count: 1155

Publisher Name: M2 Communications

Industry Names: BUSN (Any type of business); INTL (Business, International)

4/8/97 (Item 9 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

03851657 Supplier Number: 48377297 (USE FORMAT 7 FOR FULLTEXT)

SUN MICROSYSTEMS: Sun Microsystems delivers Java technology roadmap

March 25 , 1998

Word Count: 1458

Publisher Name: M2 Communications

Industry Names: BUSN (Any type of business); INTL (Business, International)

4/8/98 (Item 10 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

03851654 Supplier Number: 48377294 (USE FORMAT 7 FOR FULLTEXT)

SUN MICROSYSTEMS: Sun delivers enterprise solution for simplified Java platform deployment

March 25 , 1998

Word Count: 857

Publisher Name: M2 Communications

Industry Names: BUSN (Any type of business); INTL (Business, International)

4/8/99 (Item 11 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

03847363 Supplier Number: 48365169 (USE FORMAT 7 FOR FULLTEXT)

MOTOROLA: Intelligent GSM cable solution featuring advanced data compression technology

March 19 , 1998

Word Count: 639

Publisher Name: M2 Communications

Industry Names: BUSN (Any type of business); INTL (Business, International)

4/8/100 (Item 12 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

03829588 Supplier Number: 48317364 (USE FORMAT 7 FOR FULLTEXT)

FUJITSU: Fujitsu develops single chip MPEG2 decoder LSI for DVDs

Feb 26 , 1998

Word Count: 508

Publisher Name: M2 Communications

Industry Names: BUSN (Any type of business); INTL (Business, International)

4/8/101 (Item 13 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

03761892 Supplier Number: 48141278 (USE FORMAT 7 FOR FULLTEXT)

Windows CE: Hitachi Announces Next Generation Handheld PC With Microsoft Windows CE 2.0

Nov 24 , 1997

Word Count: 928

Publisher Name: EDGE Publishing

Company Names: *Hitachi Home Electronics (America) Inc.

Event Names: *336 (Product introduction)

Geographic Names: *1USA (United States)

Product Names: *3573160 (Personal Digital Assistants)

Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation); TELC (Telecommunications)

NAICS Codes: 334111 (Electronic Computer Manufacturing)

4/8/102 (Item 14 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

03324298 Supplier Number: 46833873 (USE FORMAT 7 FOR FULLTEXT)

Chips: C-Cube Delivers Real-Time Digital Video Encoding to Consumer PC Applications with Debut of Low-Cost Encoder Chip Family; C-Cube's CLM41xx

Oct 28 , 1996

Word Count: 1052

Publisher Name: EDGE Publishing

Company Names: *C-Cube Microsystems Inc.

Event Names: *330 (Product information)

Geographic Names: *1USA (United States)

Product Names: *3573299 (Miscellaneous Computer Peripherals NEC)

Industry Names: BUSN (Any type of business); TELC (Telecommunications)

NAICS Codes: 334119 (Other Computer Peripheral Equipment Manufacturing)

Ticker Symbols: CUBE

4/8/103 (Item 15 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

03315837 Supplier Number: 46813908 (USE FORMAT 7 FOR FULLTEXT)

-C-CUBE MICROSYSTEMS: Real-time digital video encoding for PC apps with low-cost encoder chips

Oct 21 , 1996

Word Count: 1181

Publisher Name: M2 Communications

Industry Names: BUSN (Any type of business); INTL (Business, International)

4/8/104 (Item 16 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

03028674 Supplier Number: 46186523 (USE FORMAT 7 FOR FULLTEXT)

SPECIAL REPORT: Loughborough Sound Images plc

March 1 , 1996

Word Count: 2467

Publisher Name: Architecture Technology Corporation

Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation)

4/8/105 (Item 17 from file: 636)
Gale Group Newsletter DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02809949 Supplier Number: 45700572 (USE FORMAT 7 FOR FULLTEXT)

Interactive Home's Monthly News Digest
August , 1995
Word Count: 1253
Publisher Name: Jupiter Communications
Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation)

4/8/106 (Item 18 from file: 636)
Gale Group Newsletter DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02085369 Supplier Number: 43843305 (USE FORMAT 7 FOR FULLTEXT)

CHIPS: MOTOROLA UNVEILS NEXT-GENERATION 8-BIT MICROCONTROLLER ARCHITECTURE
May 17 , 1993
Word Count: 1323
Publisher Name: EDGE Publishing
Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation); TELC
(Telecommunications)

4/8/107 (Item 1 from file: 813)
PR Newswire
(c) 1999 PR Newswire Association Inc. All rights reserved.
1075801 LATU034a
Toshiba Announces Industry's Most Complete Reference Design Tools for DVD PC Applications

Date: April 1, 1997
Word Count: 653

Company Name: TOSHIBA AMERICA ELECTRONIC COMPONENTS, INC.
Product: COMPUTER, ELECTRONICS (CPR)
Descriptors: NEW PRODUCTS & SERVICES (PDT)
State: CALIFORNIA (CA)
Section Heading: BUSINESS; TECHNOLOGY

4/8/108 (Item 2 from file: 813)
PR Newswire

(c) 1999 PR Newswire Association Inc. All rights reserved.

0528983

MN002

CRAY RESEARCH REVEALS KEY FEATURES OF FIRST MPP SYSTEM

Date: October 26, 1992

Word Count: 1,834

Company Name: CRAY RESEARCH, INC.

Ticker Symbol: CYR (NYS)

Product: COMPUTER, ELECTRONICS (CPR)

Descriptors: NEW PRODUCTS & SERVICES (PDT)

State: MINNESOTA (MN)

Section Heading: BUSINESS; TECHNOLOGY

4/8/109 (Item 1 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

06410443 Supplier Number: 54875836 (USE FORMAT 7 FOR FULLTEXT)

Data Capture Grows Wider -- Small Computing Devices And Embedded Systems Can Feed Large Data Warehouses, Leading To Potentially Powerful Data Analysis.(Technology Information)

June 14 , 1999

Word Count: 2186

Publisher Name: CMP Publications, Inc.

Event Names: *600 (Market information - general)

Geographic Names: *1USA (United States)

Product Names: *7372422 (DBMS Utilities); 7372425 (Data Warehousing Software); 7372522 (Data Acquisition Software); 7375000 (Database Providers)

Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation); TELC (Telecommunications)

NAICS Codes: 51121 (Software Publishers); 514191 (On-Line Information Services)

Special Features: LOB

4/8/110 (Item 2 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

06093672 Supplier Number: 53638382 (USE FORMAT 7 FOR FULLTEXT)

MPEG-4 systems need specialized CPUs.(Technology Information)

Jan 25 , 1999

Word Count: 975

Publisher Name: CMP Publications, Inc.

Event Names: *330 (Product information)

Geographic Names: *1USA (United States)

Product Names: *3674000 (Semiconductor Devices)

Industry Names: BUSN (Any type of business); ELEC (Electronics); ENG (Engineering and Manufacturing)

NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)

Special Features: LOB

4/8/111 (Item 3 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

06072249 Supplier Number: 53549695 (USE FORMAT 7 FOR FULLTEXT)

Will New Architectures Provide A Path To Recovery? -- The trail is still tortuous for memory-module suppliers, which are grappling with new design and test challenges.(Industry Trend or Event)

Jan 11 , 1999

Word Count: 2166

Publisher Name: CMP Publications, Inc.

Company Names: *Rambus Inc.

Event Names: *010 (Forecasts, trends, outlooks); 600 (Market information - general)

Geographic Names: *1USA (United States)

Product Names: *3573221 (Computer RAM)

Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation); ELEC (Electronics)

NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)

Special Features: COMPANY

4/8/112 (Item 4 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

05695816 Supplier Number: 50134095 (USE FORMAT 7 FOR FULLTEXT)

Memory debacle at core of economic woes -- Silver lining seen in DRAM storm cloud

July 6 , 1998

Word Count: 1357

Publisher Name: CMP Publications, Inc.

Event Names: *600 (Market information - general)

Geographic Names: *1USA (United States)

Product Names: *3674125 (Random Access Memory Circuits)

Industry Names: BUSN (Any type of business); ELEC (Electronics); ENG (Engineering and Manufacturing)

NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)

4/8/113 (Item 5 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

04777385 Supplier Number: 47032518 (USE FORMAT 7 FOR FULLTEXT)

Set-top-box design needs reassessment

Jan 13 , 1997

Word Count: 1657

Publisher Name: CMP Publications, Inc.

Event Names: *350 (Product standards, safety, & recalls)

Geographic Names: *1USA (United States)

Product Names: *3662255 (Cable Television Converters ex Addressable)

Industry Names: BUSN (Any type of business); ELEC (Electronics); ENG (Engineering and Manufacturing)

NAICS Codes: 33422 (Radio and Television Broadcasting and Wireless Communications Equipment Manufacturing)

4/8/114 (Item 6 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

04347350 Supplier Number: 46376051 (USE FORMAT 7 FOR FULLTEXT)

New RAM Burns Rubber

May 10 , 1996

Word Count: 409

Publisher Name: Boucher Communications, Inc.

Event Names: *330 (Product information); 600 (Market information - general)

Geographic Names: *1USA (United States)

Product Names: *3674125 (Random Access Memory Circuits)

Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation)

NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)

4/8/115 (Item 7 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

04153901 Supplier Number: 46065066 (USE FORMAT 7 FOR FULLTEXT)

Getting more out of Ethernet

Jan 15 , 1996

Word Count: 1872

Publisher Name: CMP Publications, Inc.

Event Names: *390 (Nonmanufacturing technology)

Geographic Names: *1USA (United States)

Product Names: *3661250 (Data Communications Systems)

Industry Names: BUSN (Any type of business); ELEC (Electronics); ENG (Engineering and Manufacturing)

NAICS Codes: 33421 (Telephone Apparatus Manufacturing)

4/8/116 (Item 8 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

03904442 Supplier Number: 45628520 (USE FORMAT 7 FOR FULLTEXT)

Philips Cooking Up Full Menu

June 26 , 1995

Word Count: 1774

Publisher Name: Cahners Publishing Company

Company Names: *Philips Semiconductors Intl

Event Names: *220 (Strategy & planning)

Geographic Names: *4EUNE (Netherlands)

Product Names: *3674000 (Semiconductor Devices)

Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation); ELEC (Electronics)

NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)

Special Features: LOB; COMPANY

4/8/117 (Item 9 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

02969597 Supplier Number: 44023522 (USE FORMAT 7 FOR FULLTEXT)

Intel, ATI in race against VESA bus

August 9 , 1993

Word Count: 1514

Publisher Name: CMP Publications, Inc.

Company Names: *ATI Technologies Inc.; Intel Corp.

Event Names: *350 (Product standards, safety, & recalls)

Geographic Names: *1USA (United States); 1CANA (Canada)

Product Names: *3573259 (Computer Output Devices NEC)

Industry Names: BUSN (Any type of business); ELEC (Electronics); ENG (Engineering and Manufacturing)

NAICS Codes: 334119 (Other Computer Peripheral Equipment Manufacturing)

Ticker Symbols: INTC

Special Features: COMPANY

4/8/118 (Item 10 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

02494360 Supplier Number: 43295852 (USE FORMAT 7 FOR FULLTEXT)

DESPITE SIMILARITIES, THERE ARE BIG DIFFERENCES BETWEEN VESA VL AND INTEL PCI:

Two local-bus standards battle it out

Sept 14 , 1992

Word Count: 1355

Publisher Name: CMP Publications, Inc.

Event Names: *460 (Use of materials & supplies); 350 (Product standards, safety, & recalls)

Geographic Names: *1USA (United States)

Product Names: *3573120 (Microcomputers); 3573291 (Computer Peripheral Interfaces)

Industry Names: BUSN (Any type of business); ELEC (Electronics); ENG (Engineering and Manufacturing)

NAICS Codes: 334111 (Electronic Computer Manufacturing); 334119 (Other Computer Peripheral Equipment Manufacturing)

4/8/119 (Item 11 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

02393542 Supplier Number: 43146799 (USE FORMAT 7 FOR FULLTEXT)

Data Topics: Officials ...

July 13 , 1992

Word Count: 308

Publisher Name: Cahners Publishing Company

Company Names: *Cray Research Inc.; Digital Equipment Corp.; Microunity Systems Engineering Inc.

Event Names: *330 (Product information); 460 (Use of materials & supplies)

Geographic Names: *1USA (United States)

Product Names: *3573111 (Supercomputers); 3674124 (Microprocessor Chips)

Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation); ELEC (Electronics)

NAICS Codes: 334111 (Electronic Computer Manufacturing); 334413 (Semiconductor and Related Device Manufacturing)

Ticker Symbols: CYR; DEC

Special Features: COMPANY

4/8/120 (Item 12 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

02371846 Supplier Number: 43114573 (USE FORMAT 7 FOR FULLTEXT)

Vendors reply to frame relay 20 Questions

July , 1992

Word Count: 956

Publisher Name: Nelson Publishing

Company Names: *American Tel & Tel; Cascade Telephone; Digital Equipment Corp.; Digitech (US); Hewlett-Packard Co.; Motorola Codex Corp.; Netrix Corp.; Northern Telecom Ltd.

Event Names: *330 (Product information); 360 (Services information)

Geographic Names: *1CANA (Canada); 1USA (United States)

Product Names: *3661255 (Packet Switches); 4811700 (Microwave Communications Services ex Satellite); 3661205 (Local Area Networks); 7372620 (Network Software)

Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation); TELC (Telecommunications)

NAICS Codes: 33421 (Telephone Apparatus Manufacturing); 513322 (Cellular and Other Wireless Telecommunications); 51121 (Software Publishers)

Ticker Symbols: DEC; HWP; NTRX; NT

Special Features: COMPANY

4/8/121 (Item 13 from file: 16)
Gale Group PROMT(R)
(c) 2008 Gale/Cengage. All rights reserved.
01560556 Supplier Number: 41911330 (USE FORMAT 7 FOR FULLTEXT)

Transceiver tester bows: Hp gear targets digital cellular equipment
March 4 , 1991
Word Count: 500
Publisher Name: CMP Publications, Inc.
Company Names: *Hewlett-Packard Co.
Event Names: *330 (Product information)
Geographic Names: *1USA (United States)
Product Names: *3825243 (Communications Test Equip)
Industry Names: BUSN (Any type of business); ELEC (Electronics); ENG (Engineering and Manufacturing)
NAICS Codes: 334515 (Instrument Manufacturing for Measuring and Testing Electricity and Electrical Signals)
Ticker Symbols: HWP
Special Features: COMPANY

4/8/122 (Item 14 from file: 16)
Gale Group PROMT(R)
(c) 2008 Gale/Cengage. All rights reserved.
01140327 Supplier Number: 41291687 (USE FORMAT 7 FOR FULLTEXT)

NASA gets memory-based net
April 23 , 1990
Word Count: 1411
Publisher Name: CMP Publications, Inc.
Company Names: *Computer Sciences Corp.; Design Analysis Associates Inc.
Event Names: *360 (Services information)
Geographic Names: *1USA (United States)
Product Names: *3661205 (Local Area Networks)
Industry Names: BUSN (Any type of business); ELEC (Electronics); ENG (Engineering and Manufacturing)
NAICS Codes: 33421 (Telephone Apparatus Manufacturing)
Ticker Symbols: CSC
Special Features: COMPANY

4/8/123 (Item 1 from file: 148)
Gale Group Trade & Industry DB
(c) 2008 Gale/Cengage. All rights reserved.
11668310 Supplier Number: 58614883 (USE FORMAT 7 OR 9 FOR FULL TEXT)
Last gasp of the graphics dinosaurs?(3dfx)(Hardware Review)(Evaluation)

Dec 23 , 1999

Word Count: 716 Line Count: 00059
Company Names: 3Dfx Interactive Inc.--Products
Industry Codes/Names: BUSN Any type of business; ENG Engineering and Manufacturing
Descriptors: Graphics boards/cards--Evaluation
Product/Industry Names: 3573293 (Computer Graphics, Sound and Video Processors)
NAICS Codes: 334119 Other Computer Peripheral Equipment Manufacturing
Trade Names: 3Dfx VSA-100 (Graphics accelerator/display board)--Evaluation
File Segment: TI File 148

4/8/124 (Item 2 from file: 148)
Gale Group Trade & Industry DB
(c) 2008 Gale/Cengage. All rights reserved.
11419246 Supplier Number: 56176661 (USE FORMAT 7 OR 9 FOR FULL TEXT)
DDR-SDRAM, high-speed, source-synchronous interfaces create design challenges.

Sept 2 , 1999
Word Count: 2575 Line Count: 00314
Industry Codes/Names: BUSN Any type of business; ENG Engineering and Manufacturing
Descriptors: Computer network equipment industry--Products; Synchronous communications--Equipment and supplies; Computer interfaces-- Innovations
Geographic Codes: 1USA United States
Product/Industry Names: 3661257 (LAN/WAN Adapters)
NAICS Codes: 33421 Telephone Apparatus Manufacturing
File Segment: TI File 148

4/8/125 (Item 3 from file: 148)
Gale Group Trade & Industry DB
(c) 2008 Gale/Cengage. All rights reserved.
10509650 Supplier Number: 53056232 (USE FORMAT 7 OR 9 FOR FULL TEXT)
Personal ATMs: Secure, Portable, Electronic Commerce With SmartPhones And Smart Cards.

Oct 1 , 1998
Word Count: 713 Line Count: 00062
Industry Codes/Names: BUSN Any type of business; CMPT Computers and Office Automation; ELEC Electronics
Descriptors: Smart cards--Usage; Electronic commerce--Forecasts; Mobile communication systems--Usage
Product/Industry Names: 3679120 (Magnetic Cards); 3662165 (Mobile Telephones ex Cellular); 3662166 (Cellular Telephones)
Product/Industry Names: 3679 Electronic components, not elsewhere classified; 3663 Radio & TV communications equipment
File Segment: TI File 148

4/8/126 (Item 4 from file: 148)
Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

10357540 Supplier Number: 20976748 (USE FORMAT 7 OR 9 FOR FULL TEXT)

DRAM market: It's beautiful to buyers.

July 16 , 1998

Word Count: 2981 Line Count: 00237

Industry Codes/Names: BUSN Any type of business; TRAN Transportation, Distribution and Purchasing

Descriptors: Dynamic random access memory--Prices and rates; Semiconductor industry-- Prices and rates

Product/Industry Names: 3573221 (Computer RAM)

Product/Industry Names: 3674 Semiconductors and related devices

File Segment: TI File 148

4/8/127 (Item 5 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

10335585 Supplier Number: 20936184 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Java For All Platforms -- New variants make it easy to develop server and mobile application components.(the maturing of the Java computing environment) (Industry Trend or Event)

July 20 , 1998

Word Count: 1474 Line Count: 00123

Industry Codes/Names: BUSN Any type of business; CMPT Computers and Office Automation; TELC Telecommunications

Descriptors: Program development software--Marketing; Java (Computer program language)--Marketing

Product/Industry Names: 7372510 (Software Development Tools)

Product/Industry Names: 7372 Prepackaged software

File Segment: CD File 275

4/8/128 (Item 6 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

10314738 Supplier Number: 20895545 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Memory debacle at core of economic woes -- Silver lining seen in DRAM storm cloud.(major suppliers announce new densities) (Industry Trend or Event)

July 6 , 1998

Word Count: 1434 Line Count: 00111

Industry Codes/Names: BUSN Any type of business; ELEC Electronics; ENG Engineering and Manufacturing

Descriptors: Semiconductor industry--Finance; Random access memory--Innovations

Product/Industry Names: 3674125 (Random Access Memory Circuits)

Product/Industry Names: 3674 Semiconductors and related devices

File Segment: CD File 275

4/8/129 (Item 7 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

09213920 Supplier Number: 19039934 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Set-top-box design needs reassessment. (Hitachi design process using Super H microprocessor illustrates need for advanced design tool in set-top box design)(Special Report on Embedded Systems; Part I: Processor Architectures) (Product Information)

Jan 13 , 1997

Word Count: 1762 Line Count: 00139

Special Features: illustration; chart

Company Names: Hitachi Ltd.--Products

Industry Codes/Names: ELEC Electronics; ENG Engineering and Manufacturing; BUSN Any type of business

Descriptors: Embedded systems--Case studies; Microprocessors--Usage; Semiconductor industry--Products

Product/Industry Names: 3674124 (Microprocessor Chips)

Product/Industry Names: 3674 Semiconductors and related devices

Trade Names: Hitachi SH (Microprocessor)--Usage

File Segment: CD File 275

4/8/130 (Item 8 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

09049373 Supplier Number: 18789615 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Upgrades: the best for the buck. (upgrading PCs) (includes related articles on differing types of RAM modules, whether CPU or RAM upgrades provide best performance boost and CPU upgrade issues) (Technology Information)

Nov , 1996

Word Count: 5101 Line Count: 00371

Special Features: illustration; table; graph

Descriptors: Random access memory--Usage; Upgrading--Equipment and supplies; Microprocessors--Usage

Product/Industry Names: 3674 Semiconductors and related devices; 3571 Electronic computers

File Segment: CD File 275

4/8/131 (Item 9 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

08831464 Supplier Number: 18389500 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Lower-power and faster devices tackle multimedia needs.

May 1 , 1996

Word Count: 4352 Line Count: 00341

Special Features: illustration; chart

Industry Codes/Names: CMPT Computers and Office Automation; ELEC Electronics
Descriptors: Custom integrated circuits--Conferences, meetings, seminars, etc. ; Custom Integrated Circuits Conference--1996
Product/Industry Names: 3674180 (Integrated Circuits by Function)
Product/Industry Names: 3674 Semiconductors and related devices
File Segment: TI File 148

4/8/132 (Item 10 from file: 148)
Gale Group Trade & Industry DB
(c) 2008 Gale/Cengage. All rights reserved.
08743772 Supplier Number: 18378216 (USE FORMAT 7 OR 9 FOR FULL TEXT)
Advances in PCMCIA-based data acquisition technology. (Personal Computer Memory Card Industry Association)

April 18 , 1996
Word Count: 2630 Line Count: 00216

Special Features: illustration; chart; graph
Company Names: National Instruments Corp.--Products
Industry Codes/Names: ENG Engineering and Manufacturing
Descriptors: Personal Computer Memory Card International Association--Standards; Boards/cards (Computers)--Product development; Semiconductor industry--Products
Product/Industry Names: 3573290 (Computer Auxiliary Eqp NEC); 3674180 (Integrated Circuits by Function)

Product/Industry Names: 3577 Computer peripheral equipment, not elsewhere classified; 3674 Semiconductors and related devices
Ticker Symbols: NATI
File Segment: TI File 148

4/8/133 (Item 11 from file: 148)
Gale Group Trade & Industry DB
(c) 2008 Gale/Cengage. All rights reserved.
08400451 Supplier Number: 17807885 (USE FORMAT 7 OR 9 FOR FULL TEXT)
Getting more out of Ethernet. (upgrading Ethernet networks for better data transmission speeds)(Design SuperCon '96: Communications Trends: Pulling ATM into the Mainstream) (Technology Information)

Jan 15 , 1996
Word Count: 2013 Line Count: 00156

Special Features: illustration; chart
Industry Codes/Names: ELEC Electronics; ENG Engineering and Manufacturing
Descriptors: Ethernet--Design and construction; Network architecture--Design and construction
Product/Industry Names: 4811250 (Local Area Networks)
Product/Industry Names: 4822 Telegraph & other communications

File Segment: CD File 275

4/8/134 (Item 12 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

08222068 Supplier Number: 17645409 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Taligent keeps its promises. (Taligent's CommonPoint Applications System object-oriented development system) (Software Review)(Evaluation)

Oct 23 , 1995

Word Count: 3277 Line Count: 00286

Special Features: illustration; table; chart

Company Names: Taligent Inc.--Products

Industry Codes/Names: CMPT Computers and Office Automation

Descriptors: Program development software--Evaluation

Product/Industry Names: 7372510 (Computer Language Software ex Military)

Product/Industry Names: 7372 Prepackaged software

Trade Names: CommonPoint Application System (Application development software)-- Evaluation

File Segment: CD File 275

4/8/135 (Item 13 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

07671370 Supplier Number: 16660900 (USE FORMAT 7 OR 9 FOR FULL TEXT)

JIT's impact on a firm's financial statements. (just-in-time inventory system)

Wntr , 1995

Word Count: 3557 Line Count: 00292

Industry Codes/Names: CNST Construction and Materials; INTL Business, International

Descriptors: Just in time inventory systems--Evaluation; Financial statements--Analysis; Business enterprises--Finance

File Segment: TI File 148

4/8/136 (Item 14 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

07608175 Supplier Number: 16531570 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Draw workstation graphics into mainstream PCs. (includes related article)

Dec 8 , 1994

Word Count: 4388 Line Count: 00342

Special Features: illustration; chart
Industry Codes/Names: ENG Engineering and Manufacturing
Descriptors: Graphics boards/cards--Design and construction
Product/Industry Names: 3573250 (Computer Optical & Graphics Eqp)
Product/Industry Names: 3577 Computer peripheral equipment, not elsewhere classified
File Segment: TI File 148

4/8/137 (Item 15 from file: 148)
Gale Group Trade & Industry DB
(c) 2008 Gale/Cengage. All rights reserved.
05792273 Supplier Number: 11950093 (USE FORMAT 7 OR 9 FOR FULL TEXT)
Alliant introduces massively parallel supercomputer. (Alliant Computer Systems Corp.'s Campus/800)
(Product Announcement)

Feb 3 , 1992
Word Count: 517 Line Count: 00046
Company Names: Alliant Computer Systems Corp.--Product introduction
Industry Codes/Names: GOVT Government and Law; CMPT Computers and Office Automation
Descriptors: Computer industry--Corrupt practices; Supercomputers--Product introduction
Product/Industry Names: 3571 Electronic computers
Trade Names: Alliant Computer Systems Campus/800 (Supercomputer)--Product introduction
File Segment: CD File 275

4/8/138 (Item 16 from file: 148)
Gale Group Trade & Industry DB
(c) 2008 Gale/Cengage. All rights reserved.
05540545 Supplier Number: 11596173 (USE FORMAT 7 OR 9 FOR FULL TEXT)
Pianos, organs & home keyboards. (Buyers Guide)

Nov , 1991
Word Count: 11269 Line Count: 00929
Industry Codes/Names: ARTS Arts and Entertainment
Descriptors: Musical instruments industry--Directories
Product/Industry Names: 3931 Musical instruments
File Segment: TI File 148

4/8/139 (Item 17 from file: 148)
Gale Group Trade & Industry DB
(c) 2008 Gale/Cengage. All rights reserved.
04844491 Supplier Number: 08901224 (USE FORMAT 7 OR 9 FOR FULL TEXT)
Industrial line-scan inspection cuts costs, improves yield. (line-scan imaging systems for industrial use have cost and speed advantages over area-scan systems)

Sept , 1990

Word Count: 1377 Line Count: 00108

Special Features: illustration; chart

Company Names: Data Translation Inc.--Products

Industry Codes/Names: ELEC Electronics; ENG Engineering and Manufacturing

Descriptors: Imaging systems--Usage; Image processing equipment industry-- Products

Product/Industry Names: 3577 Computer peripheral equipment, not elsewhere classified

Ticker Symbols: DATX

File Segment: TI File 148

4/8/140 (Item 18 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

04780626 Supplier Number: 09257987 (USE FORMAT 7 OR 9 FOR FULL TEXT)

DSP design made easy. (digital signal processors; includes related article on DSP math)

July 26 , 1990

Word Count: 3129 Line Count: 00254

Special Features: illustration; chart

Industry Codes/Names: METL Metals, Metalworking and Machinery; ELEC Electronics; ENG Engineering and Manufacturing

Descriptors: Digital signal processing; Engineering design--Equipment and supplies

File Segment: TI File 148

4/8/141 (Item 19 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

03518665 Supplier Number: 06664141 (USE FORMAT 7 OR 9 FOR FULL TEXT)

NAB offers a groaning board of technological fare. (National Association of Broadcasters equipment exhibition)

April 25 , 1988

Word Count: 10248 Line Count: 00799

Special Features: illustration; photograph

Industry Codes/Names: ARTS Arts and Entertainment

Descriptors: National Association of Broadcasters--Exhibitions; Television broadcasting --Exhibitions; Electronics industry--Exhibitions

Product/Industry Names: 3651 Household audio and video equipment; 4833 Television broadcasting stations;

3663 Radio & TV communications equipment; 3670 Electronic Components and Accessories

File Segment: TI File 148

4/8/142 (Item 20 from file: 148)
Gale Group Trade & Industry DB
(c) 2008 Gale/Cengage. All rights reserved.
03130152 Supplier Number: 04782641 (USE FORMAT 7 OR 9 FOR FULL TEXT)
NAB in the 'big D.' (National Association of Broadcasters convention, Dallas)

March 30 , 1987
Word Count: 32530 Line Count: 02895

Special Features: illustration; table
Industry Codes/Names: ARTS Arts and Entertainment
Descriptors: National Association of Broadcasters--Conferences, meetings, seminars, etc. ; Broadcasters--
Conferences, meetings, seminars, etc.
Product/Industry Names: 4833 Television broadcasting stations; 8611 Business associations
File Segment: TI File 148

4/8/143 (Item 1 from file: 20)
Dialog Global Reporter
(c) 2008 Dialog. All rights reserved.
09928215 (USE FORMAT 7 OR 9 FOR FULLTEXT)
iMobile Suite. The new features enable an administrator to view hardware and

March 06, 2000
Word Count: 337
SIC Codes/Descriptions: 7372 (Prepackaged Software)
Naics Codes/Descriptions: 51121 (Software Publishers)

4/8/144 (Item 2 from file: 20)
Dialog Global Reporter
(c) 2008 Dialog. All rights reserved.
09252748 (USE FORMAT 7 OR 9 FOR FULLTEXT)
FUJITSU: Fujitsu develops high performance graphics display controller

January 24, 2000
Word Count: 1020
Company Names: Fujitsu Ltd
Descriptors: Facilities & Equipment; Company News; New Products & Services; Marketing
Country Names/Codes: Japan (JP)
Regions: Asia; Far East; Pacific Rim
SIC Codes/Descriptions: 3812 (Search & Navigation Equipment)
Naics Codes/Descriptions: 334511 (Search Detection & Navigation Instrument Mfg)

4/8/145 (Item 3 from file: 20)

Dialog Global Reporter
(c) 2008 Dialog. All rights reserved.
09190515 (USE FORMAT 7 OR 9 FOR FULLTEXT)
Motorola Inc. - 4th Quarter & Final Results

January 18, 2000
Word Count: 4163
Company Names: Motorola Inc
Descriptors: Sales; Marketing; Company News; Sports; General News
Country Names/Codes: United States of America (US)
Regions: Americas; North America; Pacific Rim
SIC Codes/Descriptions: 7941 (Professional Sports Clubs & Promoters); 3663 (Radio & TV Communications Equipment)
Naics Codes/Descriptions: 711211 (Sports Teams & Clubs); 33422 (Radio TV Broadcast & Wireless Communications Equipment Mfg)

4/8/146 (Item 4 from file: 20)
Dialog Global Reporter
(c) 2008 Dialog. All rights reserved.
09166246 (USE FORMAT 7 OR 9 FOR FULLTEXT)
Motorola Reports Higher Fourth-Quarter, Full-Year Sales and -2-

January 17, 2000
Word Count: 1497
Company Names: Motorola Inc; Iridium World Communications Ltd; Teledesic LLC
Descriptors: Government News; Regulation of Business; Company News; Restructuring; Strategy; Facilities & Equipment; Contracts & New Orders; Sales; Marketing
Country Names/Codes: United States of America (US)
Regions: Americas; North America; Pacific Rim
SIC Codes/Descriptions: 4812 (Radiotelephone Communications); 4813 (Telephone Communications Ex Radio); 3663 (Radio & TV Communications Equipment)
Naics Codes/Descriptions: 513322 (Cellular & Other Wireless Telecommunications); 51334 (Satellite Telecommunications); 33422 (Radio TV Broadcast & Wireless Communications Equipment Mfg)

4/8/147 (Item 5 from file: 20)
Dialog Global Reporter
(c) 2008 Dialog. All rights reserved.
01297788 (USE FORMAT 7 OR 9 FOR FULLTEXT)
Compaq Unveils Powerful Armada 7800 Notebook PC Featuring -2-

April 02, 1998
Word Count: 752
Company Names: Compaq Computer Corporation
Descriptors: New Products & Services; Equities Market
Country Names/Codes: United States of America (US)

Regions: North America

Province/State: Texas

SIC Codes/Descriptions: 3571 (Electronic Computers); 3570 (Computer & Office Equipment)

4/8/148 (Item 1 from file: 635)

Business Dateline(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

0612163 95-68477

NEC Electronics Inc. debuts high-density ASIC with CBA architecture

Publication Date: 950626

Word Count: 1,094

Dateline: Mountain View, CA, US

Company Names: NEC Electronics Inc, Mountain View, CA, US, DUNS:09-853-0603, SIC:3674,

Classification Codes: 8650 (Electrical & electronics industries); 7500 (Product planning & development)

Descriptors: Electronics industry; Integrated circuits; Product introduction; Pacific

4/8/149 (Item 2 from file: 635)

Business Dateline(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

0394532 93-45922

Motorola unveils next-generation 8-bit microcontroller architecture

Publication Date: 930510

Word Count: 1,415

Dateline: Austin, TX, US

Company Names: Motorola Inc, Roselle, IL, US, DUNS:00-132-5463, SIC:3662;3674;3651;3661, Ticker:MOT

Classification Codes: 8650 (Electrical & electronics industries); 7500 (Product planning & development)

Descriptors: Electronics industry; Computer peripherals; Product introduction; Southwest

4/8/150 (Item 1 from file: 47)

Gale Group Magazine DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

05153957 Supplier Number: 20396428 (USE FORMAT 7 OR 9 FOR FULL TEXT)

All fired up. (role of neuronal rhythms in perception)

Feb 21 , 1998

Word Count: 1992 Line Count: 00167

Special Features: photograph; illustration

Descriptors: Visual perception--Research; Neurons--Research

File Segment: MI File 47

4/8/151 (Item 2 from file: 47)

Gale Group Magazine DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

03483293 Supplier Number: 09267724 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Highly parallel computation.

Nov 30 , 1990

Word Count: 6531 Line Count: 00525

Special Features: illustration; chart; graph

Descriptors: Parallel processing--Research; Electronic data processing in research--Equipment and supplies;

Computer architecture--Research

File Segment: MI File 47

4/8/152 (Item 3 from file: 47)

Gale Group Magazine DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

02740631 Supplier Number: 04038512 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Parallel processing: fact or fancy? Parallel architectures are sprouting everywhere - but no everyone who claims to have one really does.

Dec 1 , 1985

Word Count: 4735 Line Count: 00388

Special Features: illustration; chart; table

Company Names: Ametek Inc.--Innovations

Descriptors: HyperCube (computer)--Innovations; California Institute of Technology--Research; Parallel processing--Innovations; Computer architecture--Innovations

File Segment: MI File 47

4/8/153 (Item 4 from file: 47)

Gale Group Magazine DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

02733575 Supplier Number: 03835248 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Backcast; in which it is shown that forecasting technology is much like telling fortunes: you win some and you lose some.

July 1 , 1985

Word Count: 3074 Line Count: 00250

Special Features: illustration; table; chart

Descriptors: Datamation (Periodical)--Forecasts; computer industry--Forecasts; high technology--Forecasts

SIC Codes: 3571 Electronic computers; 7374 Data processing and preparation

File Segment: MI File 47

4/8/154 (Item 5 from file: 47)

Gale Group Magazine DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

02586088 Supplier Number: 03410455 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Computer software for process control.

Sept , 1984

Word Count: 5195 Line Count: 00406

Special Features: illustration; chart

Descriptors: Process control--Computer programs; Software--Design and construction

File Segment: MI File 47

4/8/155 (Item 6 from file: 47)

Gale Group Magazine DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

02586084 Supplier Number: 03410450 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Operating systems.

Sept , 1984

Word Count: 4985 Line Count: 00380

Special Features: illustration; chart; table

Descriptors: Operating systems--Design and construction

File Segment: MI File 47

? d s

Set	Items	Description
-----	-------	-------------

S1	63914638	S PD<20000329
----	----------	---------------

S2	361599	S S1 AND (BILL??? OR INVOIC??? OR CHARG??? OR PAYMENT OR PAYMENTS OR SETTLE??? OR SETTLEMENT) AND (BROKER??? OR SYNCHRO OR SYNCHRONIZ??? OR SYNCHRONIZATION OR MEDIAT??? OR MEDIATION OR INTERMEDIAT???) AND (SERVER OR COMPUTER OR SERVERS OR COMPUTERIZ??? OR COMPUTERIZATION OR COMPUTERS OR PROCESS??? OR TERMINAL OR TERMINALS OR UNIT OR UNITS OR APPARATUS)
----	--------	--

S3	250	S S2 AND ((SYNCHRO OR SYNCHRONIZ??? OR SYNCHRONIZATION) (5N) (MEMORY OR BUFFER???))
----	-----	--

S4 155 RD (unique items)

? s s4 and (purchas??? or buy??? or shop???? or consumer or consumers or customer or customers or patron or partons or commerce or ecommerce or e-commerce)

Processing

Processing

Processing

Processing

Processing

Processing

Processing

155	S4
12333960	PURCHAS???
15139860	BUY???
8108025	SHOP????
8831579	CONSUMER
5924010	CONSUMERS
9661744	CUSTOMER
15680806	CUSTOMERS
206373	PATRON
2086	PARTONS
5647559	COMMERCE
167135	ECOMMERCE
89960	E-COMMERCE

S5 111 S S4 AND (PURCHAS??? OR BUY??? OR SHOP???? OR CONSUMER OR CONSUMERS OR CUSTOMER OR CUSTOMERS OR PATRON OR PARTONS OR COMMERCE OR ECOMMERCE OR E-COMMERCE)

? t s5/free/all

5/8/1 (Item 1 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

02332051 86066948

****USE FORMAT 7 OR 9 FOR FULL TEXT****

Bottleneck allocation methodology (BAM): an algorithm

Word Count: 4251

1999

Geographic Names: United States; US

Descriptors: Algorithms; Studies; Manufacturing resource planning; Production scheduling
Classification Codes: 9190 (CN=United States); 9130 (CN=Experimental/Theoretical); 5310 (CN=Production planning & control)
Print Media ID: 11839

5/8/2 (Item 2 from file: 15)
ABI/Inform(R)
(c) 2008 ProQuest Info&Learning. All rights reserved.
01887959 05-38951
USE FORMAT 7 OR 9 FOR FULL TEXT
SDRAM memory: DRAM and beyond
Word Count: 1541 Length: 4 Pages
Second Quarter 1999
Geographic Names: US

Descriptors: DRAM; R&D; Computer industry; Capacity; Bandwidths; Performance evaluation; Comparative analysis; Technological change
Classification Codes: 9190 (CN=United States); 5230 (CN=Computer hardware); 5400 (CN=Research & development); 8651 (CN=Computer industry)

5/8/3 (Item 3 from file: 15)
ABI/Inform(R)
(c) 2008 ProQuest Info&Learning. All rights reserved.
01838320 04-89311
USE FORMAT 7 OR 9 FOR FULL TEXT
Data capture grows wider
Word Count: 2241 Length: 5 Pages
Jun 14, 1999
Company Names:
Federal Express Corp (Duns: 05-807-0459 Ticker: FDX)
Beth Israel Deaconess Medical Center-Boston MA
Cummins Engine Co Inc (Duns: 00-641-5160 Ticker: CUM)
Brooks Brothers
Microsoft Corp (Duns: 08-146-6849 Ticker: MSFT)
Geographic Names: US

Descriptors: Data mining; Trends; Data warehouses; Portable computers; Manycompanies
Classification Codes: 9190 (CN=United States); 5240 (CN=Software & systems); 5220 (CN=Data processing management)

5/8/4 (Item 4 from file: 15)
ABI/Inform(R)
(c) 2008 ProQuest Info&Learning. All rights reserved.

01835242 04-86233

USE FORMAT 7 OR 9 FOR FULL TEXT

All-optical networks

Word Count: 5031 Length: 10 Pages

Jun 1999

Geographic Names: US

Descriptors: Fiber optic networks; Network topologies; Communications equipment ; Multiplexers; Data transmission

Classification Codes: 5250 (CN=Telecommunications systems); 9190 (CN=United States)

5/8/5 (Item 5 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01778495 04-29486

USE FORMAT 7 OR 9 FOR FULL TEXT

Always on

Word Count: 866 Length: 2 Pages

Feb 15, 1999

Company Names:

Nortel Networks

Motorola Computer Group

Geographic Names: US

Descriptors: Carriers ; Technological planning; Communications networks; Reliability

Classification Codes: 9190 (CN=United States); 8330 (CN=Broadcasting & telecommunications); 5250 (CN=Telecommunications systems); 2400 (CN=Public relations)

5/8/6 (Item 6 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01673556 03-24546

USE FORMAT 7 OR 9 FOR FULL TEXT

Java for all platforms

Word Count: 1384 Length: 3 Pages

Jul 20, 1998

Company Names:

Sun Microsystems Inc (Duns: 01-304-4532 Ticker: SUNW)

Geographic Names: US

Descriptors: Java; Systems portability; Technological change; Object oriented programming

Classification Codes: 9190 (CN=United States); 5240 (CN=Software & systems)

5/8/7 (Item 7 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01514315 01-65303

USE FORMAT 7 OR 9 FOR FULL TEXT

SLDRAM Consortium puts up fight for memory

Word Count: 543 Length: 1 Pages

Oct 6, 1997

Geographic Names: US

Descriptors: Consortia; Standards; Computer memory; Computer architecture; Competition

Classification Codes: 5230 (CN=Computer hardware); 7500 (CN=Product planning & development); 9190 (CN=United States)

5/8/8 (Item 8 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01487026 01-38014

USE FORMAT 7 OR 9 FOR FULL TEXT

Nortel's 10-GBIT/S transport platform: Delivering bandwidth to build on

Word Count: 4364 Length: 11 Pages

Jul 1997

Company Names:

Nortel Communications Inc

Geographic Names: Canada

Descriptors: Bandwidths; Multiplexers; Technological change; SONET; Product development

Classification Codes: 5250 (CN=Telecommunications systems); 8650 (CN=Electrical & electronics industries); 7500 (CN=Product planning & development); 5400 (CN=Research & development); 9172 (CN=Canada)

5/8/9 (Item 9 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01218374 98-67769

USE FORMAT 7 OR 9 FOR FULL TEXT

CMG 1995 annual conference reflects the state of the industry

Word Count: 5151 Length: 12 Pages

Jan 1996

Geographic Names: US

Descriptors: Systems management; Computer memory; Performance evaluation; Conferences

Classification Codes: 5240 (CN=Software & systems); 7300 (CN=Sales & selling); 9190 (CN=United States)

5/8/10 (Item 10 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

01089259 97-38653

USE FORMAT 7 OR 9 FOR FULL TEXT

Testing wireless

Word Count: 2928 Length: 5 Pages

Sep 1995

Geographic Names: US

Descriptors: Wireless communications; Infrastructure ; Equipment testing; Maintenance management

Classification Codes: 9190 (CN=United States); 5250 (CN=Telecommunications systems); 5130 (CN=Maintenance)

5/8/11 (Item 11 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

00975539 96-24932

USE FORMAT 7 OR 9 FOR FULL TEXT

JIT's impact on a firm's financial statements

Word Count: 3312 Length: 5 Pages

Winter 1995

Descriptors: Studies; Purchasing; Just in time; Inventory management

Classification Codes: 5120 (CN=Purchasing); 9130 (CN=Experimental/Theoretical); 5330 (CN=Inventory management)

5/8/12 (Item 12 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

00901316 95-50708

USE FORMAT 7 OR 9 FOR FULL TEXT

Lead-time models of business processes

Word Count: 6133 Length: 16 Pages

1994

Descriptors: Operations research; Production planning; Time management; Business process reengineering; Management styles; Models

Classification Codes: 2600 (CN=Management science/Operations research); 5310 (CN=Production planning & control); 2200 (CN=Managerial skills); 9130 (CN=Experimental/Theoretical)

5/8/13 (Item 13 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

00627563 92-42503

****USE FORMAT 7 OR 9 FOR FULL TEXT****

Ultracomputers: A Teraflop Before Its Time

Word Count: 13520 Length: 22 Pages

Aug 1992

Descriptors: R&D; Supercomputers; Computer industry; Product development; Parallel processing ; Processing speed; Multiprocessing

Classification Codes: 5400 (CN=Research & development); 5230 (CN=Computer hardware); 8651 (CN=Computer industry); 7500 (CN=Product planning & development)

5/8/14 (Item 14 from file: 15)

ABI/Inform(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

00623706 92-38808

****USE FORMAT 7 OR 9 FOR FULL TEXT****

Vendors Reply to Frame Relay 20 Questions

Word Count: 936 Length: 2 Pages

Jul 1992

Company Names:

Cascade Communications Corp

Cisco Systems Inc (Duns: 15-380-4570 Ticker: CSCO)

Motorola-Codex

Netrix Corp

Sync Research

Geographic Names: US

Descriptors: Packet switched networks; Manyproducts; Manycompanies; Connectivity; Data transmission; Standards; Support

Classification Codes: 5250 (CN=Telecommunications systems); 9190 (CN=United States)

5/8/15 (Item 1 from file: 9)

Business & Industry(R)

(c) 2008 Gale/Cengage. All rights reserved.

01865006 Supplier Number: 24680837 (USE FORMAT 7 OR 9 FOR FULLTEXT)

Data Capture Grows Wider -- Small Computing Devices And Embedded Systems Can Feed Large Data Warehouses, Leading To Potentially Powerful Data Analysis

June 14, 1999

Word Count: 2141

Industry Names: Applications software; Computer; Mobile communications; Personal computers; Portable

computers; Software; Telecom services; Telecommunications

Product Names: Portable computers (357165); Radiotelephone communications (481200); Database software packages (737265); Applications software packages NEC (737279)

Concept Terms: All market information; Industry forecasts; Sales; Trends; Users

Geographic Names: World (WOR)

5/8/16 (Item 2 from file: 9)

Business & Industry(R)

(c) 2008 Gale/Cengage. All rights reserved.

01593920 Supplier Number: 24317129 (USE FORMAT 7 OR 9 FOR FULLTEXT)

Silver lining seen in DRAM storm cloud

July 06, 1998

Word Count: 1401

Special Features: Table

Industry Names: Electronic components; Semiconductors

Product Names: Memory integrated circuits (367445)

Concept Terms: All market information; Industry forecasts; Market size; Sales

Geographic Names: World (WOR)

5/8/17 (Item 3 from file: 9)

Business & Industry(R)

(c) 2008 Gale/Cengage. All rights reserved.

00780751 Supplier Number: 23328070 (USE FORMAT 7 OR 9 FOR FULLTEXT)

Taligent Keeps Its Promises

October 23, 1995

Word Count: 2979

Company Names: APPLE COMPUTER INC; HEWLETT-PACKARD LTD (HEWLETT-PACKARD CO); INTERNATIONAL BUSINESS MACHINES CORP; TALIGENT INC

Industry Names: Applications software; Software

Product Names: Applications software packages (737263)

Concept Terms: All product and service information; Product introduction

Geographic Names: North America (NOAX); United States (USA)

5/8/18 (Item 1 from file: 610)

Business Wire

(c) 2008 Business Wire. All rights reserved.

00206546 20000301061B2644 (USE FORMAT 7 FOR FULLTEXT)

WideBand Corp. Begins Trading

Wednesday , March 1, 2000 18:53 EST

Word Count: 407

Product Names: CORPORATE NETWORKS; NETWORKS; COMMUNICATIONS TECHNOLOGIES;
COMPUTERS ; CORPORATE; DATA COMMUNICATIONS
Event Names: TECHNOLOGY DEVELOPMENT

5/8/19 (Item 2 from file: 610)

Business Wire

(c) 2008 Business Wire. All rights reserved.

00107707 19990922265B0166 (USE FORMAT 7 FOR FULLTEXT)

Ericsson Introduces State-of-the-Art Phone for International Business Travelers

Wednesday , September 22, 1999 09:20 EDT

Word Count: 767

Company Names: TELEFON AB LM ERICSSON; OMNIPOINT CORP; EDELMAN PUBLIC RELATIONS

Geographic Names: USA; AMERICAS; NORTH AMERICA

Product Names: ELECTRONIC MAIL; MOBILE COMMUNICATIONS; NETWORKS; PORTABLE
COMPUTERS ; RADIO COMMUNICATION; TELEPHONES; COMMUNICATIONS TECHNOLOGIES;
COMPUTERS; DATA COMMUNICATIONS; TELECOMMUNICATIONS; COMPUTER HARDWARE;
MICROCOMPUTERS

5/8/20 (Item 1 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0949984 BW0132

SUN MICROSYSTEMS : Sun Releases Beta Java 2 Platform Optimized for Solaris to Help ISVs Run
Enterprise-Class Java Applications

December 09, 1998

Byline: Business Editors & Computer Writers

Word Count: 919

5/8/21 (Item 2 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0929128 BW0119

SUN MICROSYSTEMS 6 : Built for Business - JDK 1.2 for the Solaris 7 Operating Environment; Sun's Java
Applications for the World's Strongest Operating Environment - Three Times Faster Than NT

October 27, 1998

Byline: Business Editors/High Tech Writers

Word Count: 855

5/8/22 (Item 3 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0901797 BW0360

TERA COMPUTING : Tera Computer Company Unveils Supercomputer Roadmap Providing a Future for SGI/Cray T90 Users

September 01, 1998

Byline: Business Editors/Computer Writers

Word Count: 1655

5/8/23 (Item 4 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0901792 BW0358

CQN TERA COMPUTING : Tera Computer Corrects and Replaces Previous Product Announcement, BW285, TERA-COMPUTER

September 01, 1998

Byline: Business Editors/Computer Writers

Word Count: 1681

5/8/24 (Item 5 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0885263 BW1093

ROCKWELL SEMICONDUCTOR : Rockwell Semiconductor Systems is First To Take a Single T1/E1/J1 Framer Chip to Octal Density

July 27, 1998

Byline: Business Editors and High-Tech Writers

Word Count: 1293

5/8/25 (Item 6 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0830013 BW1130

COMPAQ 2 : Compaq Introduces Ultimate Video Conferencing Kit and High-Capacity Diskette Drive for Portable PCs

April 02, 1998

Byline: Business/Technology Editors

Word Count: 1204

5/8/26 (Item 7 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0830009 BW1127

COMPAQ : Compaq Unveils Powerful Armada 7800 Notebook PC Featuring Intel's Mobile Pentium II Processor

April 02, 1998

Byline: Business/Technology Editors

Word Count: 2984

5/8/27 (Item 8 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0774312 BW1463

HITACHI HOME ELEC : Hitachi Announces Next Generation Handheld PC With Microsoft Windows CE 2.0

November 17, 1997

Byline: Business Editors & Technology Writers

Word Count: 1172

5/8/28 (Item 9 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0763830 BW1178

SMART MODULAR TECH : SMART Modular Technologies Announces High-Density Registered SDRAM

Modules for High-end Systems

October 27, 1997

Byline: Business Editors/Computer Writers

Word Count: 602

5/8/29 (Item 10 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0668277 BW1376

VTEL : VTEL Assigned Patent for Multipoint Videoconference Technology

February 03, 1997

Byline: Business Editors

Word Count: 342

5/8/30 (Item 11 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0496657 BW1028

NEC ELECTRONICS : NEC Electronics Inc. Debuts High-Density ASIC With CBA Architecture; CMOS-8LHD Ideal for High-Integration Designs

June 26, 1995

Byline: Business Editors

Word Count: 956

5/8/31 (Item 12 from file: 810)

Business Wire

(c) 1999 Business Wire . All rights reserved.

0333844 BW625

MOTOROLA : Motorola unveils next-generation 8-bit microcontroller architecture

May 10, 1993

Byline: Business Editors and Computers Writers

Word Count: 1431

5/8/32 (Item 1 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02384458

Supplier Number: 60805620 (Use Format 7 Or 9 For FULL TEXT)
Legato Updates Three Products for Windows 2000.(Product Announcement)

March 8 , 2000
Word Count: 231 Line Count: 00022
Company Names: Legato Systems Inc.--Product introduction
Geographic Codes/Names: 1USA United States
Descriptors: Backup software; Network software; Networking software product introduction
Event Codes/Names: 336 Product introduction
Product/Industry Names: 7372620 (Network Software)
SIC Codes: 7372 Prepackaged software
NAICS Codes: 51121 Software Publishers
Ticker Symbols: lgto
Trade Names: Legato NetWorker 5.7 (Backup software)--Product introduction; Legato Octopus 4.0 (Backup software)--Product introduction; Legato Cluster Enterprise 4.5.1 (Network software)--Product introduction
File Segment: CD File 275

5/8/33 (Item 2 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02377084 Supplier Number: 59624577 (Use Format 7 Or 9 For FULL TEXT)
new products.

Feb , 2000
Word Count: 4177 Line Count: 00357
File Segment: CD File 275

5/8/34 (Item 3 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02365011 Supplier Number: 58925511 (Use Format 7 Or 9 For FULL TEXT)
Motorola Reports Higher Fourth-Quarter, Full-Year Sales and Earnings.(Company Financial Information)

Jan 24 , 2000
Word Count: 3240 Line Count: 00447
Company Names: Motorola Inc.--Finance
Geographic Codes/Names: 1USA United States
Descriptors: Electronics industry; Company sales/revenue; Company earnings/profit
Event Codes/Names: 830 Sales, profits & dividends
Product/Industry Names: 3601000 (Electronics)

NAICS Codes: 3359 Other Electrical Equipment and Component Manufacturing
File Segment: CD File 275

5/8/35 (Item 4 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02277748 Supplier Number: 54090766 (Use Format 7 Or 9 For FULL TEXT)
HOOKED ON PLACEBOS.

April , 1999
Word Count: 1623 Line Count: 00130
File Segment: CD File 275

5/8/36 (Item 5 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02206808 Supplier Number: 21004107 (Use Format 7 Or 9 For FULL TEXT)
Chips: Rockwell Semiconductor Systems is First To Take a Single T1/E1/J1 Framer Chip to Octal Density.(the RS8398 chip provides a single universal octal solution for physical layer termination of multiplexed voice and data traffic)(Product Announcement)

August 3 , 1998
Word Count: 1248 Line Count: 00107
Company Names: Rockwell Semiconductor Systems--Product introduction
Descriptors: Multiplexer; Networking Hardware Product Introduction
Product/Industry Names: 3674182 (Multiplexer Circuits)
SIC Codes: 3674 Semiconductors and related devices
Trade Names: Rockwell Semiconductor RS8398 (Multiplexer)--Product introduction
File Segment: CD File 275

5/8/37 (Item 6 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02159872 Supplier Number: 20480777 (Use Format 7 Or 9 For FULL TEXT)
New Notebooks: Compaq Unveils Powerful Armada 7800 Notebook PC Featuring Intel's Mobile Pentium II Processor.(Product Announcement)

April 6 , 1998
Word Count: 2953 Line Count: 00248
Descriptors: Hardware Product Introduction; Pentium II-Based Notebook
Product/Industry Names: 3573141 (Intel-Compatible Notebook Computers)
SIC Codes: 3571 Electronic computers

Trade Names: Compaq Armada 7800 (Pentium II-based notebook)--Product introduction
File Segment: CD File 275

5/8/38 (Item 7 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

02123441 Supplier Number: 20027758 (Use Format 7 Or 9 For FULL TEXT)

Windows CE: Hitachi announces next generation Handheld PC with Microsoft Windows CE 2.0. (Hitachi Handheld PC)(Product Announcement)

Nov 24 , 1997

Word Count: 983 Line Count: 00087

Company Names: Hitachi Home Electronics (America) Inc.--Product introduction

Descriptors: Hardware Product Introduction; Personal Digital Assistant

Product/Industry Names: 3573160 (Personal Digital Assistants)

SIC Codes: 3571 Electronic computers

Trade Names: Hitachi Handheld PC (Personal digital assistant)--Product introduction

File Segment: CD File 275

5/8/39 (Item 8 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

02075514 Supplier Number: 19528096 (Use Format 7 Or 9 For FULL TEXT)

Memory overload: Making sense of RAM. (Technology Information)

June 23 , 1997

Word Count: 557 Line Count: 00046

Special Features: illustration; table

Company Names: Apple Computer Inc.--Products

Descriptors: Technology Overview; RAM; DRAM; SRAM; Microcomputer Industry

Product/Industry Names: 3674125 (Random Access Memory Circuits)

SIC Codes: 3674 Semiconductors and related devices

Ticker Symbols: AAPL

File Segment: CD File 275

5/8/40 (Item 9 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

02013777 Supplier Number: 18875632 (Use Format 7 Or 9 For FULL TEXT)

Chips: C-Cube delivers real-time digital video encoding to consumer PC applications with debut of low-cost encoder chip family; C-Cube's CLM41xx. (C-Cube Microsystems CLM4100 MPEG-1 encoders)(Product Announcement)

Oct 28 , 1996

Word Count: 1181 Line Count: 00097

Company Names: C-Cube Microsystems Inc.--Product introduction

Descriptors: Hardware Product Introduction; Video Processing Equipment

Product/Industry Names: 3573250 (Computer Optical & Graphics Equip); 3662650 (Image Processing Equip)

SIC Codes: 3577 Computer peripheral equipment, not elsewhere classified

Ticker Symbols: CUBE

Trade Names: C-Cube Microsystems CLM4100 (Video processing equipment)--Product introduction

File Segment: CD File 275

5/8/41 (Item 10 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01966984 Supplier Number: 18564817

Reporting against large databases. (the role of server-based reporting engines) (Technology Information)

August , 1996

Word Count: 2211 Line Count: 00183

Special Features: illustration; chart

Descriptors: Technology Overview; DBMS; Report Generation Software; Database Design; Data Warehousing; Client/Server Architecture

SIC Codes: 7372 Prepackaged software

File Segment: CD File 275

5/8/42 (Item 11 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01944414 Supplier Number: 18371553 (Use Format 7 Or 9 For FULL TEXT)

Pentium Classic: still the one. (Overview of evaluations 101 Pentium-based systems) (individual evaluation articles searchable under "Pentium Classic: Still the One)(includes related articles on the editors' choices, Pentium vs. Pentium Pro performance, reading the Service & Reliability boxes, Pentium PC features, benchmark test results, purchasing guidelines, price/performance index, and summary of features) (Hardware Review)(Evaluation)(Cover Story)

June 25 , 1996

Word Count: 7989 Line Count: 00596

Special Features: illustration; photograph; table; chart; graph

Company Names: Dell Computer Corp.--Products; Micron Electronics Inc.--Products

Descriptors: Hardware Multiproduct Review; Pentium-Based System

SIC Codes: 3571 Electronic computers

Ticker Symbols: DELL

Trade Names: Dell Computer Dell Dimension XPS P133c (Pentium-based system)-- Evaluation; Micron

Electronics P133 Millennia (Pentium-based system)-- Evaluation; Micron Electronics P166 Millennia (Pentium-

based system)-- Evaluation
File Segment: CD File 275

5/8/43 (Item 12 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
01938349 Supplier Number: 18295599 (Use Format 7 Or 9 For FULL TEXT)
New RAM burns rubber. (synchronous DRAM) (Technology Information)

May 10 , 1996
Word Count: 430 Line Count: 00035
Company Names: Dell Computer Corp.--Products
Descriptors: DRAM; Microcomputer Industry; Technology Overview
SIC Codes: 3571 Electronic computers
Ticker Symbols: DELL
File Segment: CD File 275

5/8/44 (Item 13 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
01902071 Supplier Number: 17946189 (Use Format 7 Or 9 For FULL TEXT)
A new memory system design for commercial and technical computing products. (HP's J/K-class memory system design) (Technology Information)

Feb , 1996
Word Count: 5557 Line Count: 00429

Special Features: illustration; chart
Descriptors: Technology Overview; System Design; Semiconductor Memory
SIC Codes: 3674 Semiconductors and related devices
File Segment: CD File 275

5/8/45 (Item 14 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
01822435 Supplier Number: 17115295 (Use Format 7 Or 9 For FULL TEXT)
Philips cooking up full menu. (Philips Semiconductors reorganizing, forming new product, manufacturing and acquisition plans)

June 26 , 1995
Word Count: 2181 Line Count: 00181

Special Features: illustration; photograph; chart
Company Names: Philips Semiconductors--Planning

Descriptors: Company Operations; Company Restructuring/Company Reorganization; Company Business And Marketing
File Segment: CD File 275

5/8/46 (Item 15 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
01819197 Supplier Number: 17365620 (Use Format 7 Or 9 For FULL TEXT)
When CPUs share.(Symmetric Multiprocessing and its hidden problems)

June 16 , 1995
Word Count: 1291 Line Count: 00105
Descriptors: CPU; Processor Architecture; Multiprocessing; Technology Information ; Technology Development
File Segment: CD File 275

5/8/47 (Item 16 from file: 275)
Gale Group Computer DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
01667427 Supplier Number: 15066256 (Use Format 7 Or 9 For FULL TEXT)
At your service. (AST Research's Manhattan SMP P/60, Compaq's ProLiant 2000 5/66 M4200A, HP's NetServer 5/60 LM, Unisys' PW2 Advantage Plus 5608, Wyse Technology's 7000i 760MP application servers) (includes five related articles on recent major developments, top-rated Compaq ProLiant 2000 system, suitability to task, price/performance ratio and benchmark testing) (Hardware Review) (Evaluation)

March 15 , 1994
Word Count: 9374 Line Count: 00733

Special Features: illustration; photograph; table; graph
Company Names: AST Research Inc.--Products; Compaq Computer Corp.--Products; Hewlett-Packard Co.--Products; Unisys Corp.--Products; Wyse Technology Inc. --Products
Descriptors: Evaluation; File Server
SIC Codes: 3571 Electronic computers; 3577 Computer peripheral equipment, not elsewhere classified; 3575 Computer terminals
Ticker Symbols: ASTA; UIS; HWP; CPQ; WYS
Trade Names: AST Research Manhattan SMP P/60 (Pentium-based system)--evaluation; Compaq ProLiant 2000 5/66 M4200A (Pentium-based system)--evaluation; HP NetServer 5/60 LM (Pentium-based system)--evaluation; Unisys PW2 Advantage Plus 5608 (Pentium-based system)--evaluation; Wyse Technology 7000i 760MP (486-based system)--evaluation
File Segment: CD File 275

5/8/48 (Item 17 from file: 275)
Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01623571 Supplier Number: 14468926 (Use Format 7 Or 9 For FULL TEXT)

NetWare 4.0 for database developers. (Novell's network operating system) (Software Review) (includes related articles on changes in version 4.0 and problems that arose when installing maintenance version 4.01) (Evaluation)

Oct , 1993

Word Count: 5608 Line Count: 00432

Special Features: illustration; table

Company Names: Novell Inc.--Products

Descriptors: Evaluation; Network Operating System

SIC Codes: 7372 Prepackaged software

Ticker Symbols: NOVL

Trade Names: NetWare 4.0 (Network operating system)--evaluation

File Segment: CD File 275

5/8/49 (Item 18 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01614381 Supplier Number: 14192541 (Use Format 7 Or 9 For FULL TEXT)

Bridging the gap between structured analysis and structured design for real-time systems. (includes related article on principles of structured analysis and design) (Technical)

August , 1993

Word Count: 4551 Line Count: 00384

Special Features: illustration; chart; table

Descriptors: Real-Time System; Structured Design Techniques; Systems Analysis; System Design; New Technique; Image Processing; Medical Diagnosis

SIC Codes: 7372 Prepackaged software

File Segment: CD File 275

5/8/50 (Item 19 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01581233 Supplier Number: 13085423 (Use Format 7 Or 9 For FULL TEXT)

Utilities put the power in PBs. (PowerBook utility collections) (Product Watch)

Jan 4 , 1993

Word Count: 821 Line Count: 00066

Company Names: Apple Computer Inc.--Products; Connectix Corp.--Products; Symantec Corp.--Products

Descriptors: Desktop Utility; Software Design; Software Selection; Computer software industry

SIC Codes: 7372 Prepackaged software; 3571 Electronic computers

Ticker Symbols: AAPL; SYMC

Trade Names: Apple Macintosh PowerBook (Notebook computer)--Computer programs; Connectix PowerBook Utilities (Operating system enhancement)-- Design and construction; Norton Essentials for PowerBook (Operating system enhancement)--Design and construction

Operating Platform: Apple Macintosh

File Segment: CD File 275

5/8/51 (Item 20 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01530956 Supplier Number: 12515725 (Use Format 7 Or 9 For FULL TEXT)

Scalability. (Ultracomputers: a Terflop Before its Time)

August , 1992

Word Count: 5648 Line Count: 00467

Special Features: illustration; chart

Descriptors: Scales; Performance Improvement; Optimization; Expandability; Processor Speed; Size; Computers; Generations of Computers; Parallel Processing; Computer industry

SIC Codes: 3571 Electronic computers

File Segment: CD File 275

5/8/52 (Item 21 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01528138 Supplier Number: 12509263 (Use Format 7 Or 9 For FULL TEXT)

Vendors reply to frame relay 20 questions. (Communications Management) (Column)

July , 1992

Word Count: 996 Line Count: 00086

Descriptors: Frame Relay; Telecommunications Services Industry; Communications Management; Purchases; Hardware Selection; Communications Equipment; Packet Switch; LAN; Data communications

SIC Codes: 4800 COMMUNICATION; 3661 Telephone and telegraph apparatus; 3660 Communications Equipment

Operating Platform: Frame Relay

File Segment: CD File 275

5/8/53 (Item 22 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01502698 Supplier Number: 11957965 (Use Format 7 Or 9 For FULL TEXT)

Understanding video displays: from CRTs to shadow masks. (cathode ray tubes) (Tech Section; includes related article on video troubleshooting) (Tutorial)

March , 1992

Word Count: 2830 Line Count: 00230

Special Features: illustration; table

Descriptors: CRT Display; Color; Tutorial; Monitors

File Segment: CD File 275

5/8/54 (Item 23 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01415150 Supplier Number: 09240412 (Use Format 7 Or 9 For FULL TEXT)

Microsoft tout multimedia PC; Windows extensions to developers.

Jan 1 , 1991

Word Count: 1835 Line Count: 00149

Company Names: Microsoft Corp.--Product introduction

Descriptors: Multimedia Technology; Application Development Software; Product Development; Applications

Programming; GUI; Market Analysis

SIC Codes: 7372 Prepackaged software

Ticker Symbols: MSFT

Operating Platform: MS Windows

File Segment: CD File 275

5/8/55 (Item 24 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

01310577 Supplier Number: 07745360 (Use Format 7 Or 9 For FULL TEXT)

Technical correspondence.

Oct , 1989

Word Count: 15663 Line Count: 01233

File Segment: CD File 275

5/8/56 (Item 1 from file: 621)

Gale Group New Prod.Annou.(R)

(c) 2008 Gale/Cengage. All rights reserved.

02282849 Supplier Number: 58611094 (USE FORMAT 7 FOR FULLTEXT)

Motorola Reports Higher Fourth-Quarter, Full-Year Sales and Earnings.

Jan 17 , 2000

Word Count: 4190

Publisher Name: Business Wire

Company Names: *Iridium L.L.C.; Motorola Inc.

Geographic Names: *1USA (United States)

Product Names: *3601000 (Electronics); 3662130 (Satellite Communications Systems)
Industry Names: BUS (Business, General); BUSN (Any type of business)
SIC Codes: 3663 (Radio & TV communications equipment); 3670 (Electronic Components and Accessories)
NAICS Codes: 3359 (Other Electrical Equipment and Component Manufacturing); 33422 (Radio and Television Broadcasting and Wireless Communications Equipment Manufacturing)
Ticker Symbols: MOT

5/8/57 (Item 2 from file: 621)
Gale Group New Prod. Annou.(R)
(c) 2008 Gale/Cengage. All rights reserved.
02118226 Supplier Number: 55148447 (USE FORMAT 7 FOR FULLTEXT)
Antex Electronics Introduces New Model BX-44 to Broadcaster Series of Digital Audio Cards.

July 14 , 1999
Word Count: 687
Publisher Name: Business Wire
Geographic Names: *1USA (United States)
Product Names: *3573293 (Computer Graphics, Sound and Video Processors)
Industry Names: BUS (Business, General); BUSN (Any type of business)
SIC Codes: 3577 (Computer peripheral equipment, not elsewhere classified)
NAICS Codes: 334119 (Other Computer Peripheral Equipment Manufacturing)

5/8/58 (Item 3 from file: 621)
Gale Group New Prod. Annou.(R)
(c) 2008 Gale/Cengage. All rights reserved.
01441361 Supplier Number: 46813394 (USE FORMAT 7 FOR FULLTEXT)
C-Cube Delivers Real-Time Digital Video Encoding to Consumer PC Applications with Introduction of Low-Cost Encoder Chip Family; C-Cube's CLM41xx Product Family Transforms Digital Video on the PC Into an Active Data Type for Internet, Desktop Video, and CD-Authoring Applications.

Oct 21 , 1996
Word Count: 1111
Publisher Name: Business Wire
Company Names: *C-Cube Microsystems Inc.
Event Names: *330 (Product information)
Geographic Names: *1USA (United States)
Product Names: *3662600 (Signal Processing Equipment)
Industry Names: BUS (Business, General); BUSN (Any type of business)
NAICS Codes: 33429 (Other Communications Equipment Manufacturing)
Ticker Symbols: CUBE

5/8/59 (Item 4 from file: 621)
Gale Group New Prod. Annou.(R)
(c) 2008 Gale/Cengage. All rights reserved.

01138217 Supplier Number: 41212126 (USE FORMAT 7 FOR FULLTEXT)
SWITCHMODE CONTROLLER RUNS FROM 12V BATTERY

March 6 , 1990

Word Count: 442

Publisher Name: Various

Company Names: *Teledyne Components Inc.

Event Names: *330 (Product information)

Geographic Names: *1USA (United States); 1U9CA (California)

Product Names: *3674156 (IC Voltage Multipliers & Regulators)

Industry Names: BUS (Business, General); BUSN (Any type of business)

NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)

Trade Names: TSC9112

5/8/60 (Item 5 from file: 621)

Gale Group New Prod.Annou.(R)

(c) 2008 Gale/Cengage. All rights reserved.

01087099 Supplier Number: 40532385 (USE FORMAT 7 FOR FULLTEXT)

NATIONAL SEMICONDUCTOR AND SGS-THOMSON MICROELECTRONICS ANNOUNCE THEIR
FIRST JOINTLY DEVELOPED ISDN PRODUCTS

Oct 4 , 1988

Word Count: 905

Publisher Name: Various

Company Names: *National Semiconductor Corp.; SGS Thomson Microelectronics S.R.L.

Event Names: *380 (Strategic alliances)

Geographic Names: *1USA (United States); 1U2NY (New York)

Product Names: *3674199 (ICs by Function NEC)

Industry Names: BUS (Business, General); BUSN (Any type of business)

NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)

Ticker Symbols: NSM

Trade Names: TP3420/ST5420; TP3076/ST5076

5/8/61 (Item 1 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

04540689 Supplier Number: 58946674 (USE FORMAT 7 FOR FULLTEXT)

Fujitsu develops high performance graphics display controller.

Jan 24 , 2000

Word Count: 1059

Publisher Name: M2 Communications Ltd.

Company Names: *Fujitsu Laboratories Ltd.

Geographic Names:

*9JAPA (Japan)

Product Names: *3573293 (Computer Graphics, Sound and Video Processors); 3674000 (Semiconductor Devices)

Industry Names: BUSN (Any type of business); INTL (Business, International)

SIC Codes: 3577 (Computer peripheral equipment, not elsewhere classified); 3674 (Semiconductors and related devices)

NAICS Codes: 334119 (Other Computer Peripheral Equipment Manufacturing); 334413 (Semiconductor and Related Device Manufacturing)

5/8/62 (Item 2 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

04438582 Supplier Number: 55814615 (USE FORMAT 7 FOR FULLTEXT)

NOTEBOOK.

Sept 20 , 1999

Word Count: 2428

Publisher Name: Warren Publishing, Inc.

Company Names: *Matsushita Electric Industrial Company Ltd.

Event Names: *443 (New capacity, new plant construction)

Geographic Names: *9JAPA (Japan)

Product Names: *3679582 (Liquid Crystal Displays); 3600000 (Electrical & Electronic Equip)

Industry Names: BUSN (Any type of business); ELEC (Electronics)

NAICS Codes: 334419 (Other Electronic Component Manufacturing); 335 (Electrical Equipment, Appliance, and Component Manufacturing)

5/8/63 (Item 3 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

04040046 Supplier Number: 53398843 (USE FORMAT 7 FOR FULLTEXT)

SUN MICROSYSTEMS: Sun releases Beta Java 2 platform optimized for Solaris.

Dec 10 , 1998

Word Count: 926

Publisher Name: M2 Communications

Industry Names: BUSN (Any type of business); INTL (Business, International)

5/8/64 (Item 4 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

04000998 Supplier Number: 53145247 (USE FORMAT 7 FOR FULLTEXT)

-SUN MICROSYSTEMS: Built for business -- JDK 1.2 for the Solaris 7 operating environment.

Oct 28 , 1998

Word Count: 864

Publisher Name: M2 Communications

Company Names: *Sun Microsystems Inc.

Geographic Names: *1USA (United States)

Product Names: *3573000 (Computers & Peripherals); 7372513 (Application Development Software)

Industry Names: BUSN (Any type of business); INTL (Business, International)

NAICS Codes: 334111 (Electronic Computer Manufacturing); 51121 (Software Publishers)

5/8/65 (Item 5 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

03935574 Supplier Number: 50213530 (USE FORMAT 7 FOR FULLTEXT)

Chips: Rockwell Semiconductor Systems is First To Take a Single T1/E1/J1 Frammer Chip to Octal Density
August 3 , 1998

Word Count: 1163

Publisher Name: EDGE Publishing

Company Names: *Rockwell Semiconductor Systems

Event Names: *336 (Product introduction)

Geographic Names: *1USA (United States)

Product Names: *3674124 (Microprocessor Chips)

Industry Names: BUSN (Any type of business); TELC (Telecommunications)

NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)

5/8/66 (Item 6 from file: 636)

Gale Group Newsletter DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

03861927 Supplier Number: 48408402 (USE FORMAT 7 FOR FULLTEXT)

New Notebooks: Compaq Unveils Powerful Armada 7800 Notebook PC Featuring Intel's Mobile Pentium II
Processor

April 6 , 1998

Word Count: 2770

Publisher Name: EDGE Publishing

Company Names: *Compaq Computer Corp.

Event Names: *330 (Product information)

Geographic Names: *1USA (United States)

Product Names: *3573140 (Notebook Computers)

Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation); TELC
(Telecommunications)

NAICS Codes: 334111 (Electronic Computer Manufacturing)

Ticker Symbols: CPQ

5/8/67 (Item 7 from file: 636)
Gale Group Newsletter DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
03859812 Supplier Number: 48401089 (USE FORMAT 7 FOR FULLTEXT)

-COMPAQ: Compaq unveils powerful Armada 7800 Notebook PC featuring Intel's Mobile Pentium II Processor
April 3 , 1998
Word Count: 2286
Publisher Name: M2 Communications
Industry Names: BUSN (Any type of business); INTL (Business, International)

5/8/68 (Item 8 from file: 636)
Gale Group Newsletter DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
03859811 Supplier Number: 48401088 (USE FORMAT 7 FOR FULLTEXT)

-COMPAQ: Compaq introduces ultimate videoconferencing kit & high-capacity diskette drive for portable PCs
April 3 , 1998
Word Count: 1155
Publisher Name: M2 Communications
Industry Names: BUSN (Any type of business); INTL (Business, International)

5/8/69 (Item 9 from file: 636)
Gale Group Newsletter DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
03851654 Supplier Number: 48377294 (USE FORMAT 7 FOR FULLTEXT)

SUN MICROSYSTEMS: Sun delivers enterprise solution for simplified Java platform deployment
March 25 , 1998
Word Count: 857
Publisher Name: M2 Communications
Industry Names: BUSN (Any type of business); INTL (Business, International)

5/8/70 (Item 10 from file: 636)
Gale Group Newsletter DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
03847363 Supplier Number: 48365169 (USE FORMAT 7 FOR FULLTEXT)

MOTOROLA: Intelligent GSM cable solution featuring advanced data compression technology
March 19 , 1998
Word Count: 639

Publisher Name: M2 Communications
Industry Names: BUSN (Any type of business); INTL (Business, International)

5/8/71 (Item 11 from file: 636)
Gale Group Newsletter DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
03829588 Supplier Number: 48317364 (USE FORMAT 7 FOR FULLTEXT)

FUJITSU: Fujitsu develops single chip MPEG2 decoder LSI for DVDs
Feb 26 , 1998
Word Count: 508
Publisher Name: M2 Communications
Industry Names: BUSN (Any type of business); INTL (Business, International)

5/8/72 (Item 12 from file: 636)
Gale Group Newsletter DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
03761892 Supplier Number: 48141278 (USE FORMAT 7 FOR FULLTEXT)

Windows CE: Hitachi Announces Next Generation Handheld PC With Microsoft Windows CE 2.0
Nov 24 , 1997
Word Count: 928
Publisher Name: EDGE Publishing
Company Names: *Hitachi Home Electronics (America) Inc.
Event Names: *336 (Product introduction)
Geographic Names: *1USA (United States)
Product Names: *3573160 (Personal Digital Assistants)
Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation); TELC (Telecommunications)
NAICS Codes: 334111 (Electronic Computer Manufacturing)

5/8/73 (Item 13 from file: 636)
Gale Group Newsletter DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
03324298 Supplier Number: 46833873 (USE FORMAT 7 FOR FULLTEXT)

Chips: C-Cube Delivers Real-Time Digital Video Encoding to Consumer PC Applications with Debut of Low-Cost Encoder Chip Family; C-Cube's CLM41xx
Oct 28 , 1996
Word Count: 1052
Publisher Name: EDGE Publishing
Company Names: *C-Cube Microsystems Inc.
Event Names: *330 (Product information)
Geographic Names: *1USA (United States)

Product Names: *3573299 (Miscellaneous Computer Peripherals NEC)
Industry Names: BUSN (Any type of business); TELC (Telecommunications)
NAICS Codes: 334119 (Other Computer Peripheral Equipment Manufacturing)
Ticker Symbols: CUBE

5/8/74 (Item 14 from file: 636)
Gale Group Newsletter DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
03315837 Supplier Number: 46813908 (USE FORMAT 7 FOR FULLTEXT)

-C-CUBE MICROSYSTEMS: Real-time digital video encoding for PC apps with low-cost encoder chips
Oct 21 , 1996
Word Count: 1181
Publisher Name: M2 Communications
Industry Names: BUSN (Any type of business); INTL (Business, International)

5/8/75 (Item 15 from file: 636)
Gale Group Newsletter DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
03028674 Supplier Number: 46186523 (USE FORMAT 7 FOR FULLTEXT)

SPECIAL REPORT: Loughborough Sound Images plc
March 1 , 1996
Word Count: 2467
Publisher Name: Architecture Technology Corporation
Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation)

5/8/76 (Item 16 from file: 636)
Gale Group Newsletter DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02809949 Supplier Number: 45700572 (USE FORMAT 7 FOR FULLTEXT)

Interactive Home's Monthly News Digest
August , 1995
Word Count: 1253
Publisher Name: Jupiter Communications
Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation)

5/8/77 (Item 17 from file: 636)
Gale Group Newsletter DB(TM)
(c) 2008 Gale/Cengage. All rights reserved.
02085369 Supplier Number: 43843305 (USE FORMAT 7 FOR FULLTEXT)

CHIPS: MOTOROLA UNVEILS NEXT-GENERATION 8-BIT MICROCONTROLLER ARCHITECTURE

May 17 , 1993

Word Count: 1323

Publisher Name: EDGE Publishing

Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation); TELC (Telecommunications)

5/8/78 (Item 1 from file: 813)

PR Newswire

(c) 1999 PR Newswire Association Inc. All rights reserved.

1075801 LATU034a

Toshiba Announces Industry's Most Complete Reference Design Tools for DVD PC Applications

Date: April 1, 1997

Word Count: 653

Company Name: TOSHIBA AMERICA ELECTRONIC COMPONENTS, INC.

Product: COMPUTER, ELECTRONICS (CPR)

Descriptors: NEW PRODUCTS & SERVICES (PDT)

State: CALIFORNIA (CA)

Section Heading: BUSINESS; TECHNOLOGY

5/8/79 (Item 2 from file: 813)

PR Newswire

(c) 1999 PR Newswire Association Inc. All rights reserved.

0528983 MN002

CRAY RESEARCH REVEALS KEY FEATURES OF FIRST MPP SYSTEM

Date: October 26, 1992

Word Count: 1,834

Company Name: CRAY RESEARCH, INC.

Ticker Symbol: CYR (NYS)

Product: COMPUTER, ELECTRONICS (CPR)

Descriptors: NEW PRODUCTS & SERVICES (PDT)

State: MINNESOTA (MN)

Section Heading: BUSINESS; TECHNOLOGY

5/8/80 (Item 1 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

06410443 Supplier Number: 54875836 (USE FORMAT 7 FOR FULLTEXT)

Data Capture Grows Wider -- Small Computing Devices And Embedded Systems Can Feed Large Data Warehouses, Leading To Potentially Powerful Data Analysis.(Technology Information)

June 14 , 1999

Word Count: 2186

Publisher Name: CMP Publications, Inc.

Event Names: *600 (Market information - general)

Geographic Names: *1USA (United States)

Product Names: *7372422 (DBMS Utilities); 7372425 (Data Warehousing Software); 7372522 (Data Acquisition Software); 7375000 (Database Providers)

Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation); TELC (Telecommunications)

NAICS Codes: 51121 (Software Publishers); 514191 (On-Line Information Services)

Special Features: LOB

5/8/81 (Item 2 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

06072249 Supplier Number: 53549695 (USE FORMAT 7 FOR FULLTEXT)

Will New Architectures Provide A Path To Recovery? -- The trail is still tortuous for memory-module suppliers, which are grappling with new design and test challenges.(Industry Trend or Event)

Jan 11 , 1999

Word Count: 2166

Publisher Name: CMP Publications, Inc.

Company Names: *Rambus Inc.

Event Names: *010 (Forecasts, trends, outlooks); 600 (Market information - general)

Geographic Names: *1USA (United States)

Product Names: *3573221 (Computer RAM)

Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation); ELEC (Electronics)

NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)

Special Features: COMPANY

5/8/82 (Item 3 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

05695816 Supplier Number: 50134095 (USE FORMAT 7 FOR FULLTEXT)

Memory debacle at core of economic woes -- Silver lining seen in DRAM storm cloud

July 6 , 1998

Word Count: 1357

Publisher Name: CMP Publications, Inc.

Event Names: *600 (Market information - general)

Geographic Names: *1USA (United States)

Product Names: *3674125 (Random Access Memory Circuits)

Industry Names: BUSN (Any type of business); ELEC (Electronics); ENG (Engineering and Manufacturing)

NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)

5/8/83 (Item 4 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

04777385 Supplier Number: 47032518 (USE FORMAT 7 FOR FULLTEXT)

Set-top-box design needs reassessment

Jan 13 , 1997

Word Count: 1657

Publisher Name: CMP Publications, Inc.

Event Names: *350 (Product standards, safety, & recalls)

Geographic Names: *1USA (United States)

Product Names: *3662255 (Cable Television Converters ex Addressable)

Industry Names: BUSN (Any type of business); ELEC (Electronics); ENG (Engineering and Manufacturing)

NAICS Codes: 33422 (Radio and Television Broadcasting and Wireless Communications Equipment Manufacturing)

5/8/84 (Item 5 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

04347350 Supplier Number: 46376051 (USE FORMAT 7 FOR FULLTEXT)

New RAM Burns Rubber

May 10 , 1996

Word Count: 409

Publisher Name: Boucher Communications, Inc.

Event Names: *330 (Product information); 600 (Market information - general)

Geographic Names: *1USA (United States)

Product Names: *3674125 (Random Access Memory Circuits)

Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation)

NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)

5/8/85 (Item 6 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

03904442 Supplier Number: 45628520 (USE FORMAT 7 FOR FULLTEXT)

Philips Cooking Up Full Menu

June 26 , 1995

Word Count: 1774

Publisher Name: Cahners Publishing Company

Company Names: *Philips Semiconductors Intl

Event Names: *220 (Strategy & planning)

Geographic Names: *4EUNE (Netherlands)

Product Names: *3674000 (Semiconductor Devices)

Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation); ELEC (Electronics)

NAICS Codes: 334413 (Semiconductor and Related Device Manufacturing)

Special Features: LOB; COMPANY

5/8/86 (Item 7 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

02969597 Supplier Number: 44023522 (USE FORMAT 7 FOR FULLTEXT)

Intel, ATI in race against VESA bus

August 9 , 1993

Word Count: 1514

Publisher Name: CMP Publications, Inc.

Company Names: *ATI Technologies Inc.; Intel Corp.

Event Names: *350 (Product standards, safety, & recalls)

Geographic Names: *1USA (United States); 1CANA (Canada)

Product Names: *3573259 (Computer Output Devices NEC)

Industry Names: BUSN (Any type of business); ELEC (Electronics); ENG (Engineering and Manufacturing)

NAICS Codes: 334119 (Other Computer Peripheral Equipment Manufacturing)

Ticker Symbols: INTC

Special Features: COMPANY

5/8/87 (Item 8 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

02494360 Supplier Number: 43295852 (USE FORMAT 7 FOR FULLTEXT)

DESPITE SIMILARITIES, THERE ARE BIG DIFFERENCES BETWEEN VESA VL AND INTEL PCI:

Two local-bus standards battle it out

Sept 14 , 1992

Word Count: 1355

Publisher Name: CMP Publications, Inc.

Event Names: *460 (Use of materials & supplies); 350 (Product standards, safety, & recalls)

Geographic Names: *1USA (United States)

Product Names: *3573120 (Microcomputers); 3573291 (Computer Peripheral Interfaces)

Industry Names: BUSN (Any type of business); ELEC (Electronics); ENG (Engineering and Manufacturing)

NAICS Codes: 334111 (Electronic Computer Manufacturing); 334119 (Other Computer Peripheral Equipment)

Manufacturing)

5/8/88 (Item 9 from file: 16)

Gale Group PROMT(R)

(c) 2008 Gale/Cengage. All rights reserved.

02371846 Supplier Number: 43114573 (USE FORMAT 7 FOR FULLTEXT)

Vendors reply to frame relay 20 Questions

July , 1992

Word Count: 956

Publisher Name: Nelson Publishing

Company Names: *American Tel & Tel; Cascade Telephone; Digital Equipment Corp.; Digitech (US); Hewlett-Packard Co.; Motorola Codex Corp.; Netrix Corp.; Northern Telecom Ltd.

Event Names: *330 (Product information); 360 (Services information)

Geographic Names: *1CANA (Canada); 1USA (United States)

Product Names: *3661255 (Packet Switches); 4811700 (Microwave Communications Services ex Satellite); 3661205 (Local Area Networks); 7372620 (Network Software)

Industry Names: BUSN (Any type of business); CMPT (Computers and Office Automation); TELC (Telecommunications)

NAICS Codes: 33421 (Telephone Apparatus Manufacturing); 513322 (Cellular and Other Wireless Telecommunications); 51121 (Software Publishers)

Ticker Symbols: DEC; HWP; NTRX; NT

Special Features: COMPANY

5/8/89 (Item 1 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

11668310 Supplier Number: 58614883 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Last gasp of the graphics dinosaurs?(3dfx)(Hardware Review)(Evaluation)

Dec 23 , 1999

Word Count: 716 Line Count: 00059

Company Names: 3Dfx Interactive Inc.--Products

Industry Codes/Names: BUSN Any type of business; ENG Engineering and Manufacturing

Descriptors: Graphics boards/cards--Evaluation

Product/Industry Names: 3573293 (Computer Graphics, Sound and Video Processors)

NAICS Codes: 334119 Other Computer Peripheral Equipment Manufacturing

Trade Names: 3Dfx VSA-100 (Graphics accelerator/display board)--Evaluation

File Segment: TI File 148

5/8/90 (Item 2 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

10509650 Supplier Number: 53056232 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Personal ATMs: Secure, Portable, Electronic Commerce With SmartPhones And Smart Cards.

Oct 1 , 1998

Word Count: 713 Line Count: 00062

Industry Codes/Names: BUSN Any type of business; CMPT Computers and Office Automation; ELEC Electronics

Descriptors: Smart cards--Usage; Electronic commerce--Forecasts; Mobile communication systems--Usage

Product/Industry Names: 3679120 (Magnetic Cards); 3662165 (Mobile Telephones ex Cellular); 3662166 (Cellular Telephones)

Product/Industry Names: 3679 Electronic components, not elsewhere classified; 3663 Radio & TV communications equipment

File Segment: TI File 148

5/8/91 (Item 3 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

10357540 Supplier Number: 20976748 (USE FORMAT 7 OR 9 FOR FULL TEXT)

DRAM market: It's beautiful to buyers.

July 16 , 1998

Word Count: 2981 Line Count: 00237

Industry Codes/Names: BUSN Any type of business; TRAN Transportation, Distribution and Purchasing

Descriptors: Dynamic random access memory--Prices and rates; Semiconductor industry-- Prices and rates

Product/Industry Names: 3573221 (Computer RAM)

Product/Industry Names: 3674 Semiconductors and related devices

File Segment: TI File 148

5/8/92 (Item 4 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

10335585 Supplier Number: 20936184 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Java For All Platforms -- New variants make it easy to develop server and mobile application components.(the maturing of the Java computing environment) (Industry Trend or Event)

July 20 , 1998

Word Count: 1474 Line Count: 00123

Industry Codes/Names: BUSN Any type of business; CMPT Computers and Office Automation; TELC Telecommunications

Descriptors: Program development software--Marketing; Java (Computer program language)--Marketing

Product/Industry Names: 7372510 (Software Development Tools)

Product/Industry Names: 7372 Prepackaged software

File Segment: CD File 275

5/8/93 (Item 5 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

10314738 Supplier Number: 20895545 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Memory debacle at core of economic woes -- Silver lining seen in DRAM storm cloud.(major suppliers announce new densities) (Industry Trend or Event)

July 6 , 1998

Word Count: 1434 Line Count: 00111

Industry Codes/Names: BUSN Any type of business; ELEC Electronics; ENG Engineering and Manufacturing

Descriptors: Semiconductor industry--Finance; Random access memory--Innovations

Product/Industry Names: 3674125 (Random Access Memory Circuits)

Product/Industry Names: 3674 Semiconductors and related devices

File Segment: CD File 275

5/8/94 (Item 6 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

09213920 Supplier Number: 19039934 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Set-top-box design needs reassessment. (Hitachi design process using Super H microprocessor illustrates need for advanced design tool in set-top box design)(Special Report on Embedded Systems; Part I: Processor Architectures) (Product Information)

Jan 13 , 1997

Word Count: 1762 Line Count: 00139

Special Features: illustration; chart

Company Names: Hitachi Ltd.--Products

Industry Codes/Names: ELEC Electronics; ENG Engineering and Manufacturing; BUSN Any type of business

Descriptors: Embedded systems--Case studies; Microprocessors--Usage; Semiconductor industry--Products

Product/Industry Names: 3674124 (Microprocessor Chips)

Product/Industry Names: 3674 Semiconductors and related devices

Trade Names: Hitachi SH (Microprocessor)--Usage

File Segment: CD File 275

5/8/95 (Item 7 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

09049373 Supplier Number: 18789615 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Upgrades: the best for the buck. (upgrading PCs) (includes related articles on differing types of RAM modules, whether CPU or RAM upgrades provide best performance boost and CPU upgrade issues) (Technology Information)

Nov , 1996

Word Count: 5101 Line Count: 00371

Special Features: illustration; table; graph

Descriptors: Random access memory--Usage; Upgrading--Equipment and supplies; Microprocessors--Usage

Product/Industry Names: 3674 Semiconductors and related devices; 3571 Electronic computers

File Segment: CD File 275

5/8/96 (Item 8 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

08831464 Supplier Number: 18389500 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Lower-power and faster devices tackle multimedia needs.

May 1 , 1996

Word Count: 4352 Line Count: 00341

Special Features: illustration; chart

Industry Codes/Names: CMPT Computers and Office Automation; ELEC Electronics

Descriptors: Custom integrated circuits--Conferences, meetings, seminars, etc. ; Custom Integrated Circuits Conference--1996

Product/Industry Names: 3674180 (Integrated Circuits by Function)

Product/Industry Names: 3674 Semiconductors and related devices

File Segment: TI File 148

5/8/97 (Item 9 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

08222068 Supplier Number: 17645409 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Taligent keeps its promises. (Taligent's CommonPoint Applications System object-oriented development system) (Software Review)(Evaluation)

Oct 23 , 1995

Word Count: 3277 Line Count: 00286

Special Features: illustration; table; chart

Company Names: Taligent Inc.--Products

Industry Codes/Names: CMPT Computers and Office Automation

Descriptors: Program development software--Evaluation

Product/Industry Names: 7372510 (Computer Language Software ex Military)

Product/Industry Names: 7372 Prepackaged software

Trade Names: CommonPoint Application System (Application development software)-- Evaluation

File Segment: CD File 275

5/8/98 (Item 10 from file: 148)
Gale Group Trade & Industry DB
(c) 2008 Gale/Cengage. All rights reserved.
07671370 Supplier Number: 16660900 (USE FORMAT 7 OR 9 FOR FULL TEXT)
JIT's impact on a firm's financial statements. (just-in-time inventory system)

Wntr , 1995
Word Count: 3557 Line Count: 00292
Industry Codes/Names: CNST Construction and Materials; INTL Business, International
Descriptors: Just in time inventory systems--Evaluation; Financial statements--Analysis; Business enterprises--Finance
File Segment: TI File 148

5/8/99 (Item 11 from file: 148)
Gale Group Trade & Industry DB
(c) 2008 Gale/Cengage. All rights reserved.
05792273 Supplier Number: 11950093 (USE FORMAT 7 OR 9 FOR FULL TEXT)
Alliant introduces massively parallel supercomputer. (Alliant Computer Systems Corp.'s Campus/800)
(Product Announcement)

Feb 3 , 1992
Word Count: 517 Line Count: 00046
Company Names: Alliant Computer Systems Corp.--Product introduction
Industry Codes/Names: GOVT Government and Law; CMPT Computers and Office Automation
Descriptors: Computer industry--Corrupt practices; Supercomputers--Product introduction
Product/Industry Names: 3571 Electronic computers
Trade Names: Alliant Computer Systems Campus/800 (Supercomputer)--Product introduction
File Segment: CD File 275

5/8/100 (Item 12 from file: 148)
Gale Group Trade & Industry DB
(c) 2008 Gale/Cengage. All rights reserved.
05540545 Supplier Number: 11596173 (USE FORMAT 7 OR 9 FOR FULL TEXT)
Pianos, organs & home keyboards. (Buyers Guide)

Nov , 1991
Word Count: 11269 Line Count: 00929
Industry Codes/Names: ARTS Arts and Entertainment
Descriptors: Musical instruments industry--Directories
Product/Industry Names: 3931 Musical instruments
File Segment: TI File 148

5/8/101 (Item 13 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

04844491 Supplier Number: 08901224 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Industrial line-scan inspection cuts costs, improves yield. (line-scan imaging systems for industrial use have cost and speed advantages over area-scan systems)

Sept , 1990

Word Count: 1377 Line Count: 00108

Special Features: illustration; chart

Company Names: Data Translation Inc.--Products

Industry Codes/Names: ELEC Electronics; ENG Engineering and Manufacturing

Descriptors: Imaging systems--Usage; Image processing equipment industry-- Products

Product/Industry Names: 3577 Computer peripheral equipment, not elsewhere classified

Ticker Symbols: DATX

File Segment: TI File 148

5/8/102 (Item 14 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

03518665 Supplier Number: 06664141 (USE FORMAT 7 OR 9 FOR FULL TEXT)

NAB offers a groaning board of technological fare. (National Association of Broadcasters equipment exhibition)

April 25 , 1988

Word Count: 10248 Line Count: 00799

Special Features: illustration; photograph

Industry Codes/Names: ARTS Arts and Entertainment

Descriptors: National Association of Broadcasters--Exhibitions; Television broadcasting --Exhibitions; Electronics industry--Exhibitions

Product/Industry Names: 3651 Household audio and video equipment; 4833 Television broadcasting stations;

3663 Radio & TV communications equipment; 3670 Electronic Components and Accessories

File Segment: TI File 148

5/8/103 (Item 15 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

03130152 Supplier Number: 04782641 (USE FORMAT 7 OR 9 FOR FULL TEXT)

NAB in the 'big D.' (National Association of Broadcasters convention, Dallas)

March 30 , 1987

Word Count: 32530 Line Count: 02895

Special Features: illustration; table

Industry Codes/Names: ARTS Arts and Entertainment

Descriptors: National Association of Broadcasters--Conferences, meetings, seminars, etc. ; Broadcasters--Conferences, meetings, seminars, etc.

Product/Industry Names: 4833 Television broadcasting stations; 8611 Business associations

File Segment: TI File 148

5/8/104 (Item 1 from file: 20)

Dialog Global Reporter

(c) 2008 Dialog. All rights reserved.

09252748 (USE FORMAT 7 OR 9 FOR FULLTEXT)

FUJITSU: Fujitsu develops high performance graphics display controller

January 24, 2000

Word Count: 1020

Company Names: Fujitsu Ltd

Descriptors: Facilities & Equipment; Company News; New Products & Services; Marketing

Country Names/Codes: Japan (JP)

Regions: Asia; Far East; Pacific Rim

SIC Codes/Descriptions: 3812 (Search & Navigation Equipment)

Naics Codes/Descriptions: 334511 (Search Detection & Navigation Instrument Mfg)

5/8/105 (Item 2 from file: 20)

Dialog Global Reporter

(c) 2008 Dialog. All rights reserved.

09190515 (USE FORMAT 7 OR 9 FOR FULLTEXT)

Motorola Inc. - 4th Quarter & Final Results

January 18, 2000

Word Count: 4163

Company Names: Motorola Inc

Descriptors: Sales; Marketing; Company News; Sports; General News

Country Names/Codes: United States of America (US)

Regions: Americas; North America; Pacific Rim

SIC Codes/Descriptions: 7941 (Professional Sports Clubs & Promoters); 3663 (Radio & TV Communications Equipment)

Naics Codes/Descriptions: 711211 (Sports Teams & Clubs); 33422 (Radio TV Broadcast & Wireless Communications Equipment Mfg)

5/8/106 (Item 3 from file: 20)

Dialog Global Reporter

(c) 2008 Dialog. All rights reserved.

09166246 (USE FORMAT 7 OR 9 FOR FULLTEXT)

Motorola Reports Higher Fourth-Quarter, Full-Year Sales and -2-

January 17, 2000

Word Count: 1497

Company Names: Motorola Inc; Iridium World Communications Ltd; Teledesic LLC

Descriptors: Government News; Regulation of Business; Company News; Restructuring; Strategy; Facilities & Equipment; Contracts & New Orders; Sales; Marketing

Country Names/Codes: United States of America (US)

Regions: Americas; North America; Pacific Rim

SIC Codes/Descriptions: 4812 (Radiotelephone Communications); 4813 (Telephone Communications Ex Radio); 3663 (Radio & TV Communications Equipment)

Naics Codes/Descriptions: 513322 (Cellular & Other Wireless Telecommunications); 51334 (Satellite Telecommunications); 33422 (Radio TV Broadcast & Wireless Communications Equipment Mfg)

5/8/107 (Item 4 from file: 20)

Dialog Global Reporter

(c) 2008 Dialog. All rights reserved.

01297788 (USE FORMAT 7 OR 9 FOR FULLTEXT)

Compaq Unveils Powerful Armada 7800 Notebook PC Featuring -2-

April 02, 1998

Word Count: 752

Company Names: Compaq Computer Corporation

Descriptors: New Products & Services; Equities Market

Country Names/Codes: United States of America (US)

Regions: North America

Province/State:

Texas

SIC Codes/Descriptions: 3571 (Electronic Computers); 3570 (Computer & Office Equipment)

5/8/108 (Item 1 from file: 635)

Business Dateline(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

0612163 95-68477

NEC Electronics Inc. debuts high-density ASIC with CBA architecture

Publication Date: 950626

Word Count: 1,094

Dateline: Mountain View, CA, US

Company Names: NEC Electronics Inc, Mountain View, CA, US, DUNS:09-853-0603, SIC:3674,

Classification Codes: 8650 (Electrical & electronics industries); 7500 (Product planning & development)

Descriptors: Electronics industry; Integrated circuits; Product introduction; Pacific

5/8/109 (Item 2 from file: 635)

Business Dateline(R)

(c) 2008 ProQuest Info&Learning. All rights reserved.

0394532 93-45922

Motorola unveils next-generation 8-bit microcontroller architecture

Publication Date: 930510

Word Count: 1,415

Dateline: Austin, TX, US

Company Names: Motorola Inc, Roselle, IL, US, DUNS:00-132-5463, SIC:3662;3674;3651;3661, Ticker:MOT

Classification Codes: 8650 (Electrical & electronics industries); 7500 (Product planning & development)

Descriptors: Electronics industry; Computer peripherals; Product introduction; Southwest

5/8/110 (Item 1 from file: 47)

Gale Group Magazine DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

02740631 Supplier Number: 04038512 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Parallel processing: fact or fancy? Parallel architectures are sprouting everywhere - but no everyone who claims to have one really does.

Dec 1 , 1985

Word Count: 4735 Line Count: 00388

Special Features: illustration; chart; table

Company Names: Ametek Inc.--Innovations

Descriptors: HyperCube (computer)--Innovations; California Institute of Technology--Research; Parallel processing--Innovations; Computer architecture--Innovations

File Segment: MI File 47

5/8/111 (Item 2 from file: 47)

Gale Group Magazine DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

02733575 Supplier Number: 03835248 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Backcast; in which it is shown that forecasting technology is much like telling fortunes: you win some and you lose some.

July 1 , 1985

Word Count: 3074 Line Count: 00250

Special Features: illustration; table; chart

Descriptors: Datamation (Periodical)--Forecasts; computer industry--Forecasts; high technology--Forecasts

SIC Codes: 3571 Electronic computers; 7374 Data processing and preparation

File Segment: MI File 47

? d s

Set	Items	Description
-----	-------	-------------

S1	63914638	S PD<20000329
----	----------	---------------

S2	361599	S S1 AND (BILL??? OR INVOIC??? OR CHARG??? OR PAYMENT OR PAYMENTS OR SETTLE??? OR SETTLEMENT) AND (BROKER??? OR SYNCHRO OR SYNCHRONIZ??? OR SYNCHRONIZATION OR MEDIAT??? OR MEDIATION OR INTERMEDIAT???) AND (SERVER OR COMPUTER OR SERVERS OR COMPUTERIZ??? OR COMPUTERIZATION OR COMPUTERS OR PROCESS??? OR TERMINAL OR TERMINALS OR UNIT OR UNITS OR APPARATUS)
----	--------	--

S3	250	S S2 AND ((SYNCHRO OR SYNCHRONIZ??? OR SYNCHRONIZATION) (5N) (MEMORY OR BUFFER???))
----	-----	--

S4	155	RD (unique items)
----	-----	-------------------

S5	111	S S4 AND (PURCHAS??? OR BUY??? OR SHOP???? OR CONSUMER OR CONSUMERS OR CUSTOMER OR CUSTOMERS OR PATRON OR PARTONS OR COMMERCE OR ECOMMERCE OR E-COMMERCE)
----	-----	---

? t s5/k/90

5/K/90 (Item 2 from file: 148)

Gale Group Trade & Industry DB

(c) 2008 Gale/Cengage. All rights reserved.

Personal ATMs: Secure, Portable, Electronic Commerce With SmartPhones And Smart Cards.

Abstract: ...are, however, some constraints in the design of mobil phones that will influence how electronic commerce will be conducted. For example, because mobile phones are small, inexpensive and lightweight, they tend...

Abstract:

An estimated US\$8 billion will be spent at electronic commerce (e-commerce) Web sites this year, and that amount is projected to grow to US\$546 billion by the year 2000. The growth of e-commerce may help drive the adoption of electronic cash (e-cash) cards.

As e-cash becomes viable, people will use it to settle transactions, both at real-world retail outlets and at virtual-world Web-based e-commerce sites. E-cash card users will reload value onto their cards from the ATMs (automatic...

...operating system. Many of today's mobile phones use microcontrollers and hardware DSPs (digital signal processors) that are controlled by custom firmware written by the mobile phone manufacturer. Smart phones benefit...

...Layered and modular, the architecture provides basic OS functions (e.g., task scheduling and management, memory management, synchronization, timers); hardware and software drivers, to insulate higher-level application code from the particulars of...

...and access information.

Currently, Web site content providers have no means for collecting low-value payments (or "micropayments") for content. Billing systems aren't well-equipped to handle charging and collecting of sub-dollar payments, and sub-dollar credit card transactions aren't economically feasible. Stored-value smart cards could...

Industry Codes/Names: ...CMPT Computers and Office Automation...

Descriptors: ...Electronic commerce-- ;

19981001

? ts5/k/33

5/K/33 (Item 2 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

...Contact KL Group for a free evaluation by visiting www.klgroup.com/xrt/gauge.

Sterling Commerce's CONNECT:Enterprise

Sterling Commerce announces CONNECT:Enterprise which automates and manages the movement of information between a company's...

...available for the IBM RISC system, HP 9000 and Sun Microsystems UNIX platforms.

Contact Sterling Commerce at (800) 311-9775, or www.sterlingcommerce.com.

AppWorx, Cosort Automate Data Warehousing

AppWorx Corp. and The CoSORT Company (Innovative Routines International, Inc. -- IRI) are integrating the AppWorx process automation solution and IRI's CoSORT high-performance data sorting and manipulation product. The combined...

...loading, staging and integration and batch production tasks; while keeping tabs on resource availability and server performance. Both AppWorx and CoSORT perform high-speed loads of large volumes of data.

AppWorx...

...availability. CoSORT performs the sorting required to speed Oracle, Red Brick, DB2, Sybase and SQL Server loads, and the extraction and transformation processes that make data ready for access and

analysis.

Contact AppWorx at 877-APPWORX, or visit...

...Magnitude 3.1

OrderFusion's Orders of Magnitude 3.1 is a suite of e-commerce apps that integrate the sell-side system with Operations Resource Management (ORM) buy-side e-procurement systems. Capabilities include a personalized Web site that highlights past purchases and provides access methods to find the right products, customer-specific contract pricing and discounts, cross-sells and up-sells, restricted access to certain products...

...Legacy integration services; content management and dynamic display; a personalization and recommendation engine; and e-commerce components and integration. Total-e-Business includes a complete online store template, sonic.com, which...

...code format.

Contact Bluestone Software at (856) 727-4600, or visit www.bluestone.com.

Sterling Commerce's GENTRAN:Catalyst

Sterling Commerce introduces GENTRAN:Catalyst, its new e-business broker and the latest addition to its e-Business Process Integration (BPI) offering. GENTRAN:Catalyst's capabilities include XML translation support, routing and translation decisions...

...Enterprise for managed data exchange and provide a comprehensive E-business integration solution.

Contact Sterling Commerce at (800) 311-9775, (469) 524-2565 or www.sterlingcommerce.com.

Datametrics VisualPulse

Datametrics Systems...

...provide realtime and historical reports on network latency and packet loss for specific Web sites, servers and network nodes anywhere over the Internet or across a corporate intranet or extranet. VisualPulse...

...is running. VisualPulse can be integrated with Datametrics' VisualRoute, providing visual traceroute information to the servers or sites being monitored.

Contact Datametrics at (703) 385-7700, or visit their Web site...

...IP EXTender 4000 and Branch Office EXTender 6000, "voice/data" business operations solutions which enable customers to set up remote offices via the Internet. MCK Hardware works with POS/OE 4...

...also offers a line of Internet Traffic Management products that optimize networking efficiency through effective server load balancing.

Contact Phobos Corporation at (801) 474-9200, or visit www.phobos.com.

Visit 3M's Web site at www.3M.com/volition.

GFI Releases Eicon DIVA Server PRI Drivers

GFI's Eicon PRI drivers support Eicon Technology's new DIVA

Server PRI-30M and gives FAXmaker users the ability to use the Primary Rate Interface (PRI) telephone service. When used with FAXmaker, the DIVA Server PRI-30M offers greater bandwidth for fax transmission and enables users to send and receive...fax from any Windows application. FAXmaker for Networks/SMTP 7.0 is the first fax server to integrate with Outlook without requiring Exchange server. It is a LAN fax solution with e-mail to fax and fax to e...

...solution with Sun Microsystems' Solstice Network Client and Solstice NFS Client Products. Solstice Network Client customers who want to continue integrating their Sun Solaris applications with their PC desktops can purchase WRQ Reflection Suite for X and get WRQ's PC X server. WRQ Reflection Suite for X provides precise rendering and integration of UNIX graphical applications from...

...the data in less than 0.05 milliseconds.

An optional internal disk and battery backup unit (IDBU) is available, which allows MegaRam-370 to save data to the internal disk in...

...host-side Fibre Channel adapter fault protection for systems connected to MTI's Vivant storage servers. To do this, PathFinder corrects potential failures that may occur in the data path between the host system

and the Fibre Channel storage server; then maintains continual access to the Vivant storage server by locating and redirecting data to an alternate functioning path.

PathFinder for Windows NT is available for MTI's Vivant storage servers and is priced from \$3,500 per server.

Contact MTI at (800) 999-9MTI, (714) 970-0300, or e-mail to info@mti
...

...Workstation Solutions' Quick Restore 2.6 is available with centralized enterprise NetApp filer-to-UNIX server NDMP backup and recovery. Quick Restore 2.6 enables NetApp customers to back up to a centrally located backup server with an attached tape library. Quick Restore 2.6 features new Linux support allows customers to choose an Intel- or AMD-based Linux server as a central administrative backup server. Quick Restore lets users backup from a filer to UNIX or Linux backup server's tape library; backup from a filer to a local tape device; backup from a...

...s tape device via the network and backup from a UNIX, Linux or Windows NT server to a filer with a local tape device. Each of these capabilities comes standard and at no extra charge.

In a related announcement, Workstation Solutions announced Quick Restore 2.6 for Red Hat Linux servers and clients. Linux-based Quick Restore backup servers can manage backup and recovery in networks of Linux, Windows NT, commercial UNIX and Network...

...recovery of UNIX, Windows NT or Network Appliance data to tape libraries attached to Linux servers.

Quick Restore 2.6 will soon be available for Red Hat Linux 5 and 6. Linux single server licenses start at \$2,500; unlimited licenses for Red Hat Linux 4.2, 5 and 6 are free of charge.

Contact Workstation Solutions at (800) 487-0080, or visit www.worksta.com.

Syntax e-BizFS...SYMMETRIX subsystems; and Oracle and DB2 relational database management systems.

SST-Resource Availability can be purchased under a range of options designed to match customer requirements, with pricing beginning at \$57,000. It is available directly through Softworks or through ...

...Source-Navigator Software Developer's Kit. The Developer version is \$199. Both versions can be purchased through Cygnus' web site.

Contact Cygnus at (800) CUGNUS-1, or visit www.cygnus.com.

Filemaker Server 5

FileMaker Server 5 is a software application for hosting FileMaker Pro 5 databases. The new FileMaker Server 5 features an increased database hosting for up to 250 FileMaker Pro database clients and ...

...with volume pricing available under FileMaker volume license programs. Upgrades from previous versions of FileMaker Server are available for \$499.

Contact FileMaker at (800) 725-2747 or visit www.filemaker.com... easily integrated across Windows NT, UNIX and Linux platforms to fax-enable applications and workflow processes.

Users can fax any document from any application to automate communications. Included with VSI-FAX...

...output of various monitoring utilities into a plain English report.

It identifies performance bottlenecks, runaway processes and memory leaks; recommends changes to the tunable parameters and hardware configuration; and quantifies the...

...NetServer LXR 8500

Camintonn Corporation announced memory upgrades for HP's NetServer LXR 8500 data server PC. The new line of memory is available in modules that provide capacities that range...

...NetServer system at 256MB with a current maximum memory of 32GB.

Camintonn's high bandwidth memory synchronizes itself with the system clock that controls the CPU, allowing the SDRAM to eliminate time...

...works with in-house personnel to define load scenarios, develop test scripts, manage the testing process and provide recommendations for optimum performance. LoadRunner ActiveTest starts at \$15,000 for scheduled service...

...int.com.

Visual Numerics' JWAVE 3.0

Visual Numerics JWAVE 3.0 is a client/server solution using

Java components to develop and deploy 100 percent Java applications across an enterprise...

...IMSL C Numerical Library and first-time support for the Linux operating system. JWAVE supports servers and desktop computers running Linux, UNIX or Windows NT.

Cost for a JWAVE 3.0 starter kit, including...

...com.

APCON's AUTOSWITCH 2000

APCON's AUTOSWITCH 2000 automates switching of peripherals across multiple servers on a network. AUTOSWITCH 2000 eliminates the need to manually facilitate tape and/or drive...

...solutions include RESOLVE Enterprise Snapshot for SQL-BackTrack and RESOLVE High Speed Transaction Recovery, for customers who rely on constant access to critical information that requires precise recovery and backup. Adding...

...the holistic approach to managing e-business environments, RESOLVE for E-business Management solutions allow customers to leave their databases open while making backups and recoveries.

RESOLVE High Speed Transaction Recovery...

...BackTrack: Delivers continuous database availability and enhances database performance by reducing the impact of backup processing from hours to minutes.

Contact BMC at www.bmc.com.

WRQ and HP Sign Third...systems to a Fibre Channel-based switched fabric. McData's initial HBA products cover major

server platforms,
including Windows NT, PCI Solaris, Sun Solaris and Novell.

Contact McData at (800) 545...

20000201

? t s5/k/32

5/K/32 (Item 1 from file: 275)

Gale Group Computer DB(TM)

(c) 2008 Gale/Cengage. All rights reserved.

NetWorker's main emphasis is now on Windows 2000, according to Bill Watson, Legato's director of NT business. The multiplatform backup and

recovery software is tuned to take advantage of Windows 2000 features -- such as NTFS, meta-databases, Active Directory, certificate server, encrypted file systems -- and to be in sync with new APIs. Additionally, Legato upgraded the NetWorker 5.7 modules for Microsoft Corp.'s Exchange and SQL Server to provide online backup and restore for these applications. Legato says NetWorker will help customers bring Windows 2000 into their storage area networks.

The latest version of Octopus 4.0 data replication and protection software has new extensions, such as partial file synchronization, cache memory reallocation, and support for active/active clustering. Watson says Octopus will also support four-way clustering when Windows 2000 Datacenter Server ships.

Cluster Enterprise 4.5.1 is a release of the 4.5 version that...

...NIC as a single point of failure. Cluster Enterprise is currently entering the certification testing process for Windows 2000.

20000308

? ts5/k/17

5/K/17 (Item 3 from file: 9)

Business & Industry(R)

(c) 2008 Gale/Cengage. All rights reserved.(USE FORMAT 7 OR 9 FOR FULLTEXT)

ABSTRACT:

...development tools. This release is the first to reach Taligent's investors, which are Apple Computer, IBM and Hewlett-Packard. All three companies are transplanting CommonPoint onto their own platforms. CommonPoint...

TEXT:

...third-party development tools. The release is the first to reach Taligent's investors -- Apple Computer, IBM, and Hewlett-Packard -- which are transplanting CommonPoint onto their own platforms.

The reference release...

...the transplanters: enabling several key enterprise services, such as OLE, OpenDoc, and standard object request broker interoperability, for instance, as well as reducing run-time and disk storage demands, and

boosting...

...a hybrid system that run on Microsoft's Windows operating systems. Also, OpenStep from Next Computer Inc. specifies a cross-platform subset of the NextStep development API. Each of these products...

...C++ and Unix, neither of which is very popular in mainstream Cobol-oriented information systems shops. To exploit CommonPoint, developers must learn object design, framework design, C++, and a little Unix.

Delivering corporate information systems is difficult because customers demand more functionality, while the technologies become more complicated. In response to this challenge, the...

...toward an integrated, layered model for application design and development. With this design, we can buy more components off-the-shelf and integrate them more easily, and we can reuse more...

...me port my most sophisticated CommonPoint application from IBM AIX to Microsoft Windows NT. The process took about an hour -- including recompilation and linking -- mainly because of syntax differences between CommonPoint...

...full-featured, distributed applications that can be deployed across heterogeneous peer-to-peer and client-server environments. Conceptually, CommonPoint can be divided into two layers -- application frameworks and system services frameworks...

...BMP, and TIFF); text formats (ASCII, RTF, and Unicode); and other application file formats (word processors, graphics editors, etc.).

Printing is well supported by page and document models that are independent ...

...a hierarchy of localization resources. They can be used as a basis for a customized billing application, allowing for different line and paragraph formats, styles, and languages.

Advanced Graphics

The graphics...

...can access local or remote Oracle, Sybase, and DB2 databases for queries and modifications. Transaction processing services are provided for database concurrency control and recovery.

One important service lacking in this...

...s proprietary remote object-calling protocol. This release does not support any standard object request broker, and CommonPoint objects can't communicate with non-CommonPoint services running on different computers or in different address spaces. Some data sharing is supported through low-level, non-object...

...at run time).

Microkernel services provide an interface for managing tasks and threads, interprocess communication, synchronization, and virtual memory independently of the operating system.

CommonPoint reference release 1.0 is integrated with the Sniff...

...support for transparent relational and object data access and standard

framework interfaces to object request brokers, remote calling services, and OpenDoc and OLE, all secured by system management frameworks for authentication...

Company Names:
APPLE COMPUTER INC...

Industry Names:

? S S4 AND (mircropurchas???)

155 S4

0 MIRCROPURCHAS???

S6 0 S S4 AND (MIRCROPURCHAS???)

? b 347, 348,349

Estimated Cost Summary

Project		Client		Charge Code		Searcher		Job		Service Code	User Number
						Rob Pond				51	264751
Date		Time		SessionID		Subsession		Subaccount			
12/12/2008		11:02:27		191		3					
Data Base	Hours	Access Charge	Print Credit	Types	Prints	Report	Rank	Links	CSS	Total	
Parking	0.0002	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	
15	0.0295	2.47	0.00	0.00	0.00	0.00	0.00	0.00	0.00	2.47	
9	0.0205	1.70	0.00	0.28	0.00	0.00	0.00	0.00	0.00	1.98	
610	0.0115	0.42	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.42	
810	0.0127	0.45	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.45	
275	0.0185	1.53	0.00	1.50	0.00	0.00	0.00	0.00	0.00	3.03	

624	0.0130	1.11	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.11
621	0.0218	1.79	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.79
636	0.0328	2.73	0.00	0.00	0.00	0.00	0.00	0.00	0.00	2.73
613	0.0097	0.34	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.34
813	0.0170	0.64	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.64
16	0.0637	5.37	0.00	0.00	0.00	0.00	0.00	0.00	0.00	5.37
160	0.0193	1.62	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.62
634	0.0072	0.26	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.26
148	0.0942	8.01	0.00	0.28	0.00	0.00	0.00	0.00	0.00	8.29
20	0.1227	5.49	0.00	0.00	0.00	0.00	0.00	0.00	0.00	5.49
35	0.0020	0.07	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.07
583	0.0177	1.10	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.10
65	0.0012	0.07	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.07
2	0.0105	1.38	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.38
474	0.0192	0.72	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.72
475	0.0065	0.23	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.23
99	0.0073	0.56	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.56
256	0.0008	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
635	0.0193	1.62	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.62
570	0.0103	0.88	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.88
387	0.0035	0.11	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.11
471	0.0213	0.79	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.79
492	0.0110	0.38	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.38
494	0.0090	0.34	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.34
631	0.0107	0.38	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.38
633	0.0082	0.30	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.30

638	0.0077	0.26	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.26
640	0.0062	0.23	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.23
641	0.0072	0.26	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.26
702	0.0132	0.49	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.49
703	0.0052	0.19	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.19
704	0.0110	0.42	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.42
713	0.0085	0.30	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.30
714	0.0073	0.26	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.26
715	0.0028	0.08	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.08
725	0.0013	0.04	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.04
735	0.0070	0.26	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.26
477	0.0025	0.08	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.08
710	0.0110	0.38	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.38
711	0.0075	0.26	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.26
756	0.0020	0.08	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.08
757	0.0047	0.15	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.15
47	0.0298	2.47	0.00	0.00	0.00	0.00	0.00	0.00	0.00	2.47
Sub	0.7877	\$49.07	\$0.00	\$2.06	\$0.00	\$0.00	\$0.00	\$0.00	\$0.00	\$51.13

Totals										
Session Totals	0.8095	\$49.41		Telecom	\$12.61					\$64.08

[File 347] JAPIO Dec 1976-2008/Aug(Updated 081208)

(c) 2008 JPO & JAPIO. All rights reserved.

[File 348] EUROPEAN PATENTS 1978-200849

(c) 2008 European Patent Office. All rights reserved.

[File 349] PCT FULLTEXT 1979-2008/UB=20081204\UT=20081127

(c) 2008 WIPO/Thomson. All rights reserved.

? s pd<20000329

Processing

Processing

Processing

Processing

Processing

S1 8104706 S PD<20000329

? S S1 AND (BILL??? OR INVOIC??? OR CHARG??? OR PAYMENT OR PAYMENTS OR SETTL??? OR SETTLEMENT or cashier) AND (BROKER??? OR SYNCHRO OR SYNCHRONIZ??? OR SYNCHRONIZATION OR MEDIAT??? OR INTERMEDIAT??? or intermediary or intermediaries) AND (SERVER OR COMPUTER OR SERVERS OR COMPUTERIZ??? OR COMPUTERIZATION OR COMPUTERS OR PROCESS??? OR TERMINAL OR TERMINALS OR UNIT OR UNITS OR APPARATUS)

Processing

Processing

Processing

Processing

Processing

Processing

Processing

8104706 S1

114557 BILL???

6517 INVOIC???

950198	CHARG???
42814	PAYMENT
10194	PAYMENTS
90814	SETTL???
15527	SETTLEMENT
2581	CASHIER
8984	BROKER???
3134	SYNCHRO
181366	SYNCHRONIZ???
111777	SYNCHRONIZATION
157450	MEDIAT???
657204	INTERMEDIAT???
37922	INTERMEDIARY
2333	INTERMEDIARIES
216006	SERVER
832749	COMPUTER
65457	SERVERS
29697	COMPUTERIZ???
534	COMPUTERIZATION
154160	COMPUTERS
3777446	PROCESS???
892550	TERMINAL
305058	TERMINALS
1557383	UNIT
1000440	UNITS
1614578	APPARATUS

S2 84483 S S1 AND (BILL???

OR INVOIC???

OR CHARG???

OR PAYMENT OR PAYMENTS OR

SETTL???

OR SETTLEMENT OR CASHIER) AND (BROKER???

OR SYNCHRO OR SYNCHRONIZ???

OR SYNCHRONIZATION OR MEDIAT???

OR INTERMEDIAT???

OR INTERMEDIARY OR INTERMEDIARIES) AND

(SERVER OR COMPUTER OR SERVERS OR COMPUTERIZ???

OR COMPUTERIZATION OR COMPUTERS OR

PROCESS???

OR TERMINAL OR TERMINALS OR UNIT OR UNITS OR APPARATUS)

? S S2 AND ((SYNCHRO OR SYNCHRONIZ???

OR SYNCHRONIZATION) (5N) (MEMORY OR BUFFER???)

Processing

Processing

84483 S2

	3134	SYNCHRO
	181366	SYNCHRONIZ???
	111777	SYNCHRONIZATION
	898913	MEMORY
	512103	BUFFER???
	8798	((SYNCHRO OR SYNCHRONIZ???) OR SYNCHRONIZATION) (5N) (MEMORY OR BUFFER???)
S3	923	S S2 AND ((SYNCHRO OR SYNCHRONIZ??? OR SYNCHRONIZATION) (5N) (MEMORY OR BUFFER???))

? S S3 AND (micropurchas??? or PURCHAS??? OR BUY??? OR SHOP???? OR CONSUMER OR CONSUMERS OR CUSTOMER OR CUSTOMERS OR PATRON OR PARTONS OR COMMERCE OR ECOMMERCE OR E-COMMERCE)

	923	S3
	1	MICROPURCHAS???
	150434	PURCHAS???
	32230	BUY???
	46424	SHOP????
	102236	CONSUMER
	47484	CONSUMERS
	100426	CUSTOMER
	54407	CUSTOMERS
	2367	PATRON
	27	PARTONS
	34460	COMMERCE
	1677	ECOMMERCE
	0	E-COMMERCE
S4	192	S S3 AND (MICROPURCHAS??? OR PURCHAS??? OR BUY??? OR SHOP???? OR CONSUMER OR CONSUMERS OR CUSTOMER OR CUSTOMERS OR PATRON OR PARTONS OR COMMERCE OR ECOMMERCE OR E-COMMERCE)

? s 4 and (settlement or settle or settling or settled)

Processing

Processing

Processing

Processing

Processing

Processing

```

5925394      4
      15527    SETTLEMENT
      38308    SETTLE
      39250    SETTLING
      26411    SETTLED
S5      79833    S 4 AND (SETTLEMENT OR SETTLE OR SETTLING OR SETTLED)

?  S s4 AND (SETTLEMENT OR SETTLE OR SETTLING OR SETTLED)
      192      S4
      15527    SETTLEMENT
      38308    SETTLE
      39250    SETTLING
      26411    SETTLED
S6      50      S S4 AND (SETTLEMENT OR SETTLE OR SETTLING OR SETTLED)

```

? rd

>>>W: Duplicate detection is not supported for File 347.

Duplicate detection is not supported for File 348.

Duplicate detection is not supported for File 349.

All specified files are unsupported, command ignored.

? t s6/k/all

6/K/1 (Item 1 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Country	Number	Kind	Date
---------	--------	------	------

Abstract ...1 information elements by creating a set of N non-repetitive information codes in a computer memory and by distributing among the players using communication channels a plurality of signals containing... ...players and containing information about the bets. The method further involves generating data on the payment of the bets and on the drawing process thereof during the various rounds. The signals comprising information on the bets are recorded in...

Type	Pub. Date	Kind	Text
Available Text	Language	Update	Word Count
Total Word Count (Document A)			

Total Word Count (Document B)
Total Word Count (All Documents)

Specification: ...30 information elements of the game set, are sold to players prior to the drawing process takes place. The aforementioned gambling combination is compiled in 6 lines each containing 5 elements... ...divided into 2 groups each containing 3 lines. During TV-broadcast of the prize-drawing process, in which under computer supervision only sold tickets are taken into consideration, a game operator chooses successively and randomly... ...coincidence occurs. The winning tickets are determined during several rounds of the prize fund drawing process: the first round winner is a ticket in which all 5 elements of any line... ...rusloto.ru). In much the same way as in "Russkoje Loto", a prize fund drawing process is the basis for a known lottery game "Bingo" and its TV-versions (see, Internet... ...player are in fact predetermined by a unique gambling combination of information elements of a purchased ticket and the prize-drawing outcomes are fully dependent on actions of the game organizers... ...marks on them and a wager drawing at a predetermined TV-broadcast time in a process of random choosing a combination of winning elements from the game set, while indicating the... ...the wagers of the game players do not affect in any way the prize-drawing process, which does not promote the attractiveness of the game.

Another drawback with the "Sportloto" lottery... ...or fraud actions of some employees among the considerable number of those involved in this process.

The influence of the aforementioned factors of the organization complexity of running "Sportloto"-like lotteries... ...drawing is realized by means of lottery coupons distribution and wager registration using remote electronic terminals at lottery retail outlets, said terminals transmitting via telephone networks to a host computer signals with information on sold coupons and their wager marks (US Patent No. 5,186... ...Lotto Million" with weekly prize-drawings through a national Russian TV channel.

To simplify the process of lottery running and, in particular to reduce a volume paper documents processed, a computerized telephone game system has been invented on the basis of a touch-tone telephone set and a host computer, which by means of ramified algorithm and a set of various prerecorded voice messages provides callers with: registration as game players identified by personal identification numbers (PINs), game credits purchase via credit cards, wager enter and registration of virtual game coupons associated with the wagers... ...1995, A63F 9/22). This technical solution also provides for the integration of the host computer with a caller's phone number identification unit and a billing software for charging the identified phone numbers, which makes the system accessible to players without credit cards. In... ...selected from predetermined sets of numbers as wagers until the beginning of the prize-drawing process by means of random choosing the winning combinations by the game organizer; in so doing... ...part (6 or more digits) coincide with the winning sequence; in so doing, the wager charging procedure is similar to that used for charging phone number holders for telecommunication. In this game, the overall number of various combinations equaled... ...same playing interest level as "Sportloto"-type lotteries.

In comparison with lotteries using lottery coupons, computerized telephone games are not only more simple from the point of view of the game process organization but are more attractive to potential players because of accessibility of the telephony.

The... ...method implemented in the Canadian state lottery "6 of 49", which enables participation (apart from buyers of lottery coupons at remote electronic terminals) to any user of the international computer network Internet having at least one of the worldwide accepted credit cards for paying his... ...this game, the game set of $N=49$ elements is designated in a central game computer as a set of N non-repeating information codes each being the binary code of... ...to participate in the game, via the Internet telecommunication lines on his or her personal computer, receives signals carrying information about the players' registration form and possible versions of charging the wager costs. Filled-in registration forms with marked payment options are returned as a set of signals through the Internet channels to the central computer and, when verified, recorded to the long-term memory for further identification of players and charging of wager costs. Thereupon, signals carrying information about the game set elements and also

information... rounds in which the wager drawing will be run, are transmitted from the central game computer to the personal computer of a registered player. From this information, every player chooses a combination of 6 elements... and, before the next round begins, transmits via the Internet channels to the central game computer signals identifying the given player and carrying information about his or her wagers. In the host computer, signals received from the players are identified, registered and memorized while charging the players who have sent the signals. When it is time to begin the next... whom the prize fund, as part of the total game round budget collected from the payment of wagers, is allocated in accordance with the game regulations.

Despite the technical attractiveness of... a rare probability of a large win, and the complete exclusion of players from the process of forming the winning information elements. These disadvantages lower the playing attractiveness of the games... by means of generation of a set of N non-repeating information codes in a computer memory, propagation, among the players through communication lines, of signals carrying information about the game... the wagers, identification and registration of signals received through feedback lines, forming of a wager payment data, a wager drawing within playing rounds, wherein signals carrying wager information are registered as... is completed, and the wager drawing is carried out by means of an iterative-analytical process of forming a quantitative wager distribution among the game set elements, said process is kept hidden from the players until the playing round is completed, and within every iteration of the said process a regular signal of a registered signal sequence is correlated with the information code of... said conditions are observed, and in the presence of registered wager information carrying signals not processed by the iterative-analytical wager drawing process before the completion of the current playing round the said signals are processed by the iterative-analytical wager drawing process within one of the next rounds.

Distinctive features of the iterative-analytical wager drawing process which differ one particular embodiment of the presently claimed method from another and consist in the conditions of completion of the iterative-analytical wager drawing process, affords a varying degree of attractiveness and profitableness of the game for both players and... wagering game method is a method in which within every iteration of a wager drawing process, or starting with the N iteration, information codes are revealed with which no signal has been correlated within the current playing round, and the iterative-analytical wager drawing process is completed in the processing of a signal containing information about a wager on the only game set element with whose information code no signal has been correlated by the iterative-analytical process within the current playing round before the processing of this signal.

This particular embodiment of the wagering game method called "Force of Zero... completed, makes it possible to reveal a winner right after his or her wager is processed, that is from the point of view of players, it is the game with an... game method is a method according to which within every iteration of a wager drawing process, or starting with the $(2N-1)$ iteration, information codes are revealed with which only one signal has been correlated within the current playing round, and the iterative-analytical wager drawing process is completed in the presence of only one said information code and in the absence of information codes with which no signal has been correlated by the iterative-analytical process within the current playing round before the processing of this signal.

This particular embodiment of the: wagering game method called "Force of Minimum... always makes it possible to reveal a winner right after his or her wager is processed, since a wager on a winning game set element may be fixed in the quantitative... of the current playing round during several iterations before the completion of the iterative-analytical process, thereby resulting in a conclusion of players that such a game is the game with... the game which is no less than $50\% (1-0.5N)$ of the total sum charged for wagers participating in the game.

Yet another embodiment of a wagering game method with... and organizers of the game. In that embodiment, within every iteration of a wager drawing process, or starting with the $2N$ iteration, information codes are... signals have been correlated within the current playing round, and the iterative-analytical wager drawing process is completed in the presence of only one information code with which the minimum number... not always makes it possible to reveal a winner right after his or her wager is processed, since a wager on a winning game set element may be fixed in the quantitative... of the current playing round during several iterations before the completion of

the iterative-analytical process, thereby resulting in a conclusion of players that such a game is the game with...
...profitability level of the game which is no less than 50 % of the total sum

charged for wagers participating in the game.

The solution of the aforementioned task with the achievement... ..is completed, and the wager drawing is carried out by means of an iterative-analytical process of forming a quantitative wager distribution among the game set elements, said process is kept hidden from the players until the playing round is completed, and within every iteration of the said process a regular coupon of a registered coupon sequence is correlated with the game set information... ..are observed, and in the presence of registered coupons carrying wager marks which are not processed by the iterative-analytical wager drawing process before the completion of the current playing round, these coupons are processed by the iterative-analytical wager drawing process within one of the next rounds.

Particular cases of this embodiment of a wagering game... ..game method called "Force of Zero", in which within every iteration of a wager drawing process, or starting with the N iteration, the game set elements are revealed with which no... ..within the current playing round, and the wager drawing is completed with an iterative-analytical processing of a coupon containing a wager mark corresponding to the only game set element with which no coupon has been correlated within the current playing round before the processing of this coupon;

- A wagering game method called "Force of Minimum", in which within every iteration of a wager drawing process, or starting with the (2N-1) iteration, the game set elements are revealed with which... ..game method called "Force of Minimax", in which within every iteration of a wager drawing process, or starting with the 2N iteration, the game set elements are revealed with which the... ..number of coupons have been correlated within the current playing round, and the wager drawing process is completed in the presence of only one element with which the minimum number of... ..cases of the embodiment described earlier.

A degree of influence of a player on the process and outcome of a particular game is increased in case where the both embodiments of...player in exchange for a wager which is placed without his or her participation and processed by a wager drawing process out of turn.

The attractiveness of wagering games in their particular embodiments "Force of Minimum... ..wager information which were received from the said player are withdrawn from a wager drawing process in the order opposite to that of their registration.

The present invention is also embodied in a wagering game apparatus intended for the implementation of the first of the above-described embodiments of the claimed method. Referring to FIG. 1, the wagering game apparatus comprises a game set forming unit (1) connected via data dissemination unit (2) to one of inputs of a processor (3) connected with its information output to a recognition and identification unit (4), a wager payment unit (5), a wager registration unit (6), a controller (7), a playing-logic unit (8), and a recording unit (9) which are connected in series, a playing-round counter (10) connected to the second input of the wager registration unit (6) and to the second output of the controller (7) connected with its second input to the output of the game set forming unit (1), a long-term memory unit (14) interconnected with the recognition and identification unit (4) and the wager payment unit (5), a timer (17) connected to the controller (7), the recognition and identification unit (4), the wager payment unit (5), and the recording unit (9). The apparatus also comprises a wager distribution processor (11) interconnected with the controller (7), a wager registration confirmation unit (12) connected to the input of the processor (3) and the second output of the wager registration unit (6), a payment registration unit (15) and an outcome review unit (16) which are interconnected with the long-term memory unit (14) and the processor (3) and also connected to corresponding outputs of the recognition and identification unit (4), the outputs of the recording unit (9) and the wager registration confirmation unit (12) being connected to corresponding inputs of the long-term memory unit (14).

To embody the wagering game method in its first embodiment which enables the players not to wager by themselves but to entrust the apparatus with this duty, the apparatus in accordance with the present invention additionally comprises a wager generator (13) interconnected with the recognition and identification unit (4) and also connected to the output of the game set forming unit (1).

To embody the wagering game method in its first embodiment which enables the players... ..be provided with information about current quantitative wager distribution among the game set elements, the apparatus in accordance with the present invention additionally comprises a wager drawing display unit (18) coupled between an additional output of the controller (7) and an additional input of the processor (3).

To embody the wagering game method in its first embodiment which enables the players to withdraw their wagers from a wager drawing process in the games "Force of Minimum" and "Force of Minimax", the apparatus in accordance with the present invention additionally comprises a wager returning unit (19) interconnected with the controller (7) and the long-term memory unit (14) and also connected to an output of the recognition and identification unit (4) and an input of the input/output processor (3).

With reference to FIG. 2 showing the apparatus to carry out the wagering game method in its first embodiment called "Force of Zero", a wager distribution processor (11) comprises a decoder (20) connected with its outputs to driving inputs of flip-flops... ..with its output to reset inputs of the flip-flops (21).

A modification of the apparatus having such a wager distribution processor (11), while fully conforming to the game method "Force of Zero", is characterized in that... ..accumulate data about a precise quantitative wager distribution among the game set elements in the processor (11).

Additional possibilities to accumulate data about a precise quantitative wager distribution among the game set elements in the processor (11) of this particular case of implementation of the apparatus to carry out the, wagering game method in its embodiment called "Force of Zero", and... ..moment, this data to the controller (7) are solved in case where a wager distribution processor (11) interconnected with a controller (7) comprises a decoder (20) connected with its outputs to inputs of counters (23) whose outputs are connected, via comparison units (24), to inverse inputs of a "logical AND" gate (25) connected with its output to... ..inputs of the counters (23) (see FIG. 3).

With reference to FIG. 4 showing the apparatus to carry out the wagering game method in its first embodiment called "Force of Minimum", a wager distribution processor (11) interconnected with a controller (7) comprises a decoder (20), each of N 1-bit... ..decoder being coupled to a stage of a counter (23) and of a null-comparison unit (24) and a 1-comparison unit (27) which are connected in parallel to the said counter (23), a "logical AND" gate (25) with N inverse inputs each coupled to the output of the corresponding null-comparison unit (24), an "exclusive OR" gate (28) with N inputs each coupled to the output of the corresponding 1-comparison unit (27), a "logical AND" gate (29) with two inputs connected to outputs of the gates... ..encoder (30) with N inputs each coupled to the output of the corresponding 1-comparison unit (27), the said gate (29) being connected with its output to reset inputs of the encoder (30).

With reference to FIG. 5 showing the apparatus to carry out the wagering game method in its first embodiment called "Force of Minimax", a wager distribution processor (11) interconnected with a controller (7) comprises a decoder (20), each of N 1-bit... ..decoder being coupled to a stage of a counter (23) and of a null-comparison unit (24), a minimum-comparison unit (31), and a maximum-comparison unit (32) which are connected in parallel to the said counter (23), a "logical AND" gate (25) with N inverse inputs each coupled to the output of the corresponding null-comparison unit (24), a first "exclusive OR" gate (28-1) with N inputs each coupled to the output of the corresponding minimum-comparison unit (31), a first "logical AND" gate (29-1) with two inputs connected to outputs of... ..28-2) with N inputs each coupled to the output of the corresponding maximum-comparison unit (32), a second "logical AND" gate (29-2) with two inputs connected to outputs of... ..30-1) with N inputs each coupled to the output of the corresponding minimum-comparison unit (31), a second encoder (30-2) with N inputs each coupled to the output of the corresponding maximum-comparison unit (32), a minimum-counter (33) coupled to the output of

the gate (29-1), the... counter (33) being connected with its output to the input of each of minimum-comparison units (31), and a maximum-counter (34) coupled to the output of the gate (28-2... counter (34) being connected with its output to the input of each of maximum-comparison units (32), the said gate (29-2) being connected with its output to reset inputs of... had to accompanying drawings wherein:

FIG. 1 is a function circuit of a wagering game apparatus.

FIG. 2 is a function and logic circuit of a wager drawing processor (11) for a particular embodiment of an apparatus realizing a modification of a "Force of Zero" method, without accumulation in the processor (11) of absolute quantitative data about wager distribution among the game set elements.

FIG. 3 is a function and logic circuit of a wager drawing processor (11) for a particular embodiment of an apparatus realizing a modification of a "Force of Zero" method, with accumulation in the processor (11) of absolute quantitative data about wager distribution among the game set elements.

FIG. 4 is a function and logic circuit of a wager drawing processor (11) for a particular embodiment of an apparatus realizing a modification of a "Force of Minimum" method.

FIG. 5 is a function and logic circuit of a wager drawing processor (11) for a particular embodiment of an apparatus realizing a modification of a "Force of Minimax" method.

FIG. 6 shows an example of... of a quantitative wager distribution among the game set elements formed by an iterative-analytical process for a wager drawing within an unfinished playing round according to a modification of a... of a quantitative wager distribution among the game set elements formed by an iterative-analytical process for a wager drawing within a completed playing round according to a modification of a... Minimax" method.

BEST MODE OF CARRYING OUT THE INVENTION

In the following description of an apparatus to carry out the first of the above-described embodiments of a wagering game method and modifications of an apparatus, when characterizing functional units, elements and a description of their operation, use is made of highly specialized terms and... to be addressed by them. Structural elements and the description of operation of iterative-analytical processor units are set forth in more detail using terms of function and logic circuits accepted in... messages and commands, etc.) are only mentioned as needed to define more exactly functions of units when they address certain application tasks.

It should be understood that the invention is not... scope of accepted terminology, and each term or designation used encompasses all equivalent elements and units operated in a similar way and used to perform the same functions.

A function circuit of a wagering game apparatus set forth in FIG. 1 comprises a game set forming unit (1) (the said set consisting of $N > 1$ game elements) which is connected, via a data dissemination unit (2), to an input/output processor (3). An information output of the processor (3) is connected to a recognition and identification unit (4), a wager payment unit (5), a wager registration unit (6), a controller (7), a playing-logic unit (8), and a recording unit (9) which are connected in series. A playing-round counter (10) is connected to the second input of the wager registration unit (6) and to the second output of the controller (7) interconnected (here and throughout the... understood as the availability of communication lines for data exchange) with the game set forming unit (1). The second output of the wager registration unit (6) is connected, via a wager registration confirmation unit (12), to the processor (3).

An output of the game set forming unit (1) is also connected to a third... the controller (7) and to a wager generator (13) interconnected with the recognition and identification unit (4). A long-term memory unit (14) is interconnected with the following units: the recognition and identification unit (4), the wager payment unit (5), a payment

registration unit (15), an outcome review unit (16) which in turn is interconnected with the processor (3) and connected to the output of the recognition and identification unit (4), the latter being connected with its other output to an input of the payment registration unit (15) interconnected with the input/output processor (3). Outputs of the recording unit (9) and the wager registration confirmation unit (12) are connected to corresponding inputs of the long-term memory unit (14).

Synchronization of the apparatus operation is ensured by introduction of a timer (17) and its connection to units and elements requiring synchronization (not shown in FIG. 1). Moreover, an output of the timer (17) is connected to the controller (7), the recognition and identification unit (4), the wager payment unit (5), and the recording unit (9).

Provision of the players, on their request, with information about a current wager distribution among the game set elements is carried out by a wager drawing display unit (18) coupled between an additional output of the controller (7) and an additional input of the input/output processor (3).

Return to the players, on their request, of wagers from unfinished playing rounds is carried out by a wager returning unit (19) interconnected with the controller (7) and the long-term memory unit (14) and also coupled between the recognition and identification unit (4) and the input/output processor (3) (these two connections are shown in FIG. 1 with a dotted line).

The game... via a console via a program interface to the electronic memory of a game set forming unit (1) of N non-repeating information codes, each corresponding to one of the game set elements. Information codes of the game set elements are transferred to a data dissemination unit (2) where they are converted into a message format about the game set contents and invitation to wager, and the said invitation enters a processor (3).

An input/output processor (3) converts received internal signals of the apparatus into signal to be transmitted through external communication lines and transmits these signals to registration-playing terminals (not shown in the Figure) of registered and potential players through external communication channels (not shown in the Figure) which may be computer interface lines, analog and digital telephone channels, communication links of local and global computer networks and on-line services, asynchronous communication links of cable TV networks, etc. Signals received from registration-playing terminals through external communication channels are converted by the input/output processor (3) into signals suitable by types and formats for dissemination and processing by functional units of the apparatus.

At the stage of forming a game set, the processor (3) converts a message about the game set contents received from the data dissemination unit (2) into data communication signals in accordance with standards and protocols accepted in communication links and networks used, and then sends these signals through communication channels to registration-playing terminals at a preset repetition frequency and/or on their requests.

The game set forming unit (1) may be used in a form of a computer terminal with a console and main memory, for example a personal computer; the data dissemination unit (2) may be made as a dedicated main memory array of this terminal under control of a special program; the input/output processor (3), depending upon types of applied interfaces, communication channels and networks, may be used in a form of serial and parallel computer ports, network interface boards, modems and modem pools, controllers and adapters of integrated-service digital telephone networks (ISDN), controllers of asynchronous cable TV networks, specialized telecommunication processors, Internet-protocol converters - jointly with relevant hardware-software drivers.

A registration-playing terminal is used for registration of the players, input of wager data and output of drawing outcomes. These functions may be performed by touch-tone telephone sets, electronic terminals of points of sales of lottery coupons, terminals for servicing credit cards, bank's cards and also smart-cards, including automatic teller machines, discount card service terminals, video-terminals of asynchronous cable networks, personal computers of users of corporate and global computer networks and on-line services, etc. A particular case of a registration-

playing terminal is a pulse-dial telephone set whose application in conjunction with the presently disclosed apparatus will be set forth separately.

Having received signals about game set contents and game regulations from the processor (3) to the registration-playing terminal, an operator of the registration-playing terminal selects one of the following three modes of interaction with the apparatus proposed by the data dissemination unit (2): payment registration mode, outcome review mode or wager placement mode, while sending a corresponding signal through feedback channels to the processor (3). Signals received through feedback channels are transferred, following conversion of types and formats, by the input/output processor (3) to the recognition and identification unit (4) which first of all adds an indication of a system timer (17) to the format of this signal and recognizes one of the aforementioned three modes of the apparatus - registration-playing terminal interaction.

The payment registration mode providing, as a rule, a pre-registration of personal data of players and... for the purposes of safeguarding against unauthorized access, depending upon the type of registration-playing terminal used and the mode of wager payment, may be realized differently. In case where the registration-playing terminal is used in a form of municipal or corporate network or a cable network video-terminal then, under preliminary agreement between the game authority and network authority, personal data of a subscriber may be transferred from the subscribers' database to the processor (3) for registration as a player; in so doing, the mode of wager and prize payment may be selected either the same as in the case of network service payment, or using credit cards, or in accordance with a discount-bonus system (see below). In case where the registration-playing terminal is used in a form of a terminal for servicing credit, bank's or smart cards, then all player's data necessary for registration of payments may be read-out directly from the card; in case of a smart card - including... the electronic circuit of the card, provided that the same card will be used as payment means during the game. If the game organizer is a discount system authority and the registration-playing terminal is used in a form of a terminal for servicing discount cards, then registration of a player may be made on the basis of registration data of a discount program participant, while effecting

payments in accordance with a discount-bonus system. If the game organizer is a corporate computer network or on-line service authority, for example the Internet-network service provider, then registration... basis of registration data of a corporate network or on-line service user, while effecting payments for participation in the game in accordance with a discount-bonus system and/or pre-agreed scheme of service payment. If the game organizer is an

Internet-network server holder, then the players fill in a registration form transferred from the processor (3) to their personal computers, while choosing as a payment means one of credit cards acceptable to the game organizer. Finally, if the registration-playing terminal is used in a form of an electronic terminal of the point of sales of lottery coupons, then the registration data is used as... code of a coupon which takes the place of personal data of a player, whereas payment for a wager is charged directly by an operator of the registration-playing terminal, and a payment registration mode is combined with a wager distribution mode. A type of the registration-playing terminal used is determined by the processor (3) using one of the known methods in the course of connection installation.

In any case, after the recognition and identification unit (4) has recognized the payment registration mode, signals received through feedback channels from the registration-playing terminal, said signals containing address data of the players, personal codes of access to an account-register and payments data, are converted in the processor (3) from formats of communication through external lines into an internal signal format used in the apparatus, and entered a long-term memory unit (14) via a payment registration unit (15) which carries on a dialog with a particular registration-playing terminal via the processor (3) and a connection established through communication channels. In so doing, the payment registration unit (15) assigns a database personal address to every new player and, using this address, introduces personal identification data of the player with an access code in the long-term memory unit (14) and opens a personal playing account-register of the player. In the course of... balance of his or

her account-register is supplemented with an initial amount of monetary units used to measure the cost of wagers in case of conducting games with wagers of... depending upon the power of the game set N) and with non-monetary forms of payment.

One of possible non-monetary forms of payment for the participation in the game is a discount-bonus system widely used in discount programs, where for each purchase of a good or service a personal account of the program participant is supplemented with a certain amount of bonus points to be determined by the money price of the purchase; subsequently, said points may be exchanged for new goods and/or services or used to... asynchronous TV networks, corporate networks and on-line services. In all these and similar cases, payment of wagers and prizes may be measured by certain quanta of services to be provided... player may change the amount of a personal playing balance repeatedly, while choosing again the payment registration mode and interacting with his or her account-register from the registration-playing terminal through communication channels, the input/output processor (3) and the payment registration unit (15) in order to enter new payment units to the account-register or withdraw them from the account-register for use beyond the game. As a matter of fact, the structure of identification data of the player, payment mode and payment equivalent are determined by the game organizer and fixed in the registration of the player; subsequently, this data may be defined more accurately as needed in the payment registration mode.

In the outcome review mode, an outcome review unit (16) via the processor (3) and communication channels carries on a dialog with registered and potential players, while providing them with information about an outcome of completed playing rounds from the long-term memory unit (14).

If the player chooses the wager placement mode, in response to the aforementioned message of the data dissemination unit (2) containing an invitation to wager, the player inputs his or her identification data in... as a signal of certain type and format through feedback channels to the input/output processor (3). A signal received and converted by the processor (3) is recognized by the recognition and identification unit (4) as a signal containing a players' data and wager information code, thereupon the unit (4) identifies the player when interacting with the long-term memory unit (14) while finding there his or her database address and checking the state of a... of the game method is the fact that in case of failure to identify, the unit (4) transfers a received signal to the long-term memory unit (14) thereby freeing itself for the identification of the next signal, and the unit (14) is preset via the payment registration unit (15) and the processor (3) through communication channels of dialog with the registration-playing terminal as to a more precise definition of identification data of the player and/or entering of additional payments to his or her personal account-register.

In order that the players could realize a possibility not to wager by themselves but to entrust the apparatus with this duty, the player with the help of the registration-playing terminal inputs, in a wager field of the signal sent in the wager placement mode, a... inquiry code which, following a successful signal identification, is interpreted by the recognition and identification unit (4) as a wager placement code without participation of the player; as a result, the recognition and identification unit (4) sends an inquiry for a wager to a wager generation unit (13) and, after having received from the wager generation unit (13) a wager in a form of an information code of a corresponding game set... this code for a special wager inquiry code in the identified signal. The wager generation unit (13) may be realized on the basis of one of the methods of pseudo-random-integer generation known in the computer science.

In case of a successful completion of identification, a signal with a wager is transferred from the recognition and identification unit (4) to a wager payment unit (5) which, while interacting with a personal account-register of an identified player in the long-term memory unit (14), generates for the account-register a message about one wager price's write-off, while sending this message to the long-term memory unit (14) to correct a balance of the identified account-register.

Prior to the balance correction, the long-term memory unit (14) may inquire, via the payment registration unit (15), the input/output processor (3), communication channels and the registration-playing terminal, the player about an

access code to his or her account-register if such inquiry... ..account-register is needed, a signal containing a wager is also transferred from the wager payment unit (5) to the long-term memory unit (14) in order to free the wager payment unit (5) for subsequent signals with wagers, as it is provided for in the recognition and identification unit (4) in each case of unsuccessful identification. In case of unsuccessful confirmation of an access code, the payment registration unit (15) carries on a dialog with a registration-playing terminal via the processor (3) to define more accurately an access code. Following confirmation of the access code, a signal containing a wager comes back to the wager payment unit (5) where its format is supplemented with data from the system timer (17) about the current time of paying for a wager.

If the player uses as a registration-playing terminal an electronic cash terminal of the point of sales of lottery coupons, a signal entering the input/output processor (3) contains a unique alphanumeric code of identification of a lottery coupon and a wager... ..a special code of inquiry for a wager generation on behalf of the player. When processing such a signal, the recognition and identification unit (4) jointly with the long-term memory unit (14) creates for this signal a database address of the received signal, said address being... ..the lottery coupon paid by the player. Based on this address, the long-term memory unit (14) forms an account-register of the paid coupon and sends the received signal to the wager payment unit (5) in which this signal acquires the value of system time.

As a result, in all cases a signal at the output of the wager payment unit (5) contains three completed data fields: an account-register database address, a wager value in... ..of a game set element, and system time of a signal output from the wager payment unit (5).

The recognition and identification unit (4), the wager payment unit (5), the payment registration unit (15) and the outcome review unit (16) may be realized as specialized imperative software modules located in the computer memory.

The long-term memory unit (14) is a database control system and, properly speaking, a database which are implemented with the use of high-performance disk drives of required capacity.

An output of the wager payment unit (5) enters a wager registration unit (6) where its format is supplemented with two data fields: a current round number field and a current round wager number field. To fill in these fields, the wager registration unit (6) uses readings of a playing-round counter (10) and readings of its internal counter copy of a signal so registered is transferred to a wager registration confirmation unit (12) which forms from this copy a message about a wager registration and transfers this message to the long-term memory unit (14) to input registration data of a wager in an account-register using a database address indicated in the copy, and also transfers this message via the input/output

processor (3) and communication channels to the registration-playing terminal from which a registered wager was placed. Thereupon, a registered signal is truncated by format... ..to the controller (7).

In this way, a message transferred by the wager registration confirmation unit (12) to the long-term memory unit (14) and the processor (3) followed by its transfer to the registration-playing terminal contains 5 filled in data fields: a database address, a wager value, a system time... ..controller (7) serves to load wagers contained in information code signals to a wager distribution processor (11), take decisions about the start and end of playing rounds and transfer data about quantitative distribution of wagers in the completed round to a playing-logic unit (8). In so doing, the controller (7) interacts with the wager distribution processor (11) which realizes an iterative-analytical process of forming a quantitative wager distribution among the game set elements.

In case of an... ..temporary buffer storage formed in its internal memory with signals coming from the wager registration unit (6), which may be realized in practice by means of the main memory dynamic allocation... ..start of the next round, unloads in sequence this buffer storage to the playing-logic unit (8) while extending a format of signals unloaded from the buffer storage with data about... ..of a regular information code of the wager from the

controller (7), the wager distribution processor (11) carries out iteration of the code processing and updates an accumulated wager distribution while correlating this code with an information code of a corresponding game set element formed by the game set forming unit (1) and the controller (7) loaded in the game initialization into the wager distribution processor (11), hereby determining the number of wagers correlated with a given game set element in... ..buffer storage, whereas an information wager code contained in the signal - to the wager distribution processor (11). When the signal $F=1$ comes, the controller (7) unloads from the wager distribution processor (11) data about quantitative wager distribution and about revealed special points of this distribution (last... ..minimum, absolute maximum, etc.) in the round completed and transfers thereof to the playing-logic unit (8) along with readings of the system timer about the round completion time, simultaneously transferring... ..flag F indicative of the playing round completion which is generated by the wager distribution processor (11) to the controller (7) is at the same time a winning flag in respect of the last wager processed ($F=1$ - wager won, $F=0$ - wager lost). For this reason, with respect to the... ..analytical game with instantaneous outcome in accordance with the present invention, a wager registration confirmation unit (12) provides for a delay in the generation of its signal before the completion of processing by the wager distribution processor (11) of a wager which corresponds to this signal, and supplements this signal with the flag F (see, a dotted linkage between units (7) and (12) in FIG. 1) in order to ensure "instantaneousness" of the game... ..result, in case of the iterative-analytical game with instantaneous outcome, the wager registration confirmation unit (12) produces a signal about winnings or loss of a wager placed and transfers this signal to the long-term memory unit (14) and the input/output processor (3) to be passed to a player who has placed this wager. At the same... ..conserve the common character of statement of the invention, one may consider that in the apparatus for games with instantaneous outcome a variable-length signal buffer storage is formed in the controller (7).

The playing-logic unit (8) having received from the controller (7) quantitative data about a final distribution of wagers...round and a quantitative allocation of the prize fund among prize-winning wagers. Then the unit (8) allows successive passage of signals unloaded from the buffer storage of the controller (7) to the recording unit (9), while comparing information codes of wagers with codes of winning elements of the game... ..7), provided that the codes coincide.

From the received signals with winning statuses, the recording unit (9) creates a consolidated protocol-register of the completed round which contains both general data... ..serial number in the round, a winning flag, the amount of the winnings). The recording unit (9) may also be realized on the basis of a variable-length buffer storage which... ..A consolidated protocol-register of the completed round is transferred to the long-term memory unit (14) for subsequent storage and presentation via the outcome review unit (16) and the processor (3) through communication channels to those who so desire. Upon arrival of this consolidated protocol-register, the long-term memory unit (14), in accordance with this protocol, corrects accounts-registers of the players whose wagers were... ..accounts-registers of natural persons and playing coupons whose wagers were introduced from registration-cash terminals, is of no importance.

Thus completes the processing of results of the completed playing round, and a new round begins with the arrival of the first signal containing a new round number from the wager registration unit (6) to the controller (7); and just this moment is fixed by the controller (7... ..period for a new round beginning.

An embodiment of the structure and operation of the apparatus set forth above is called everywhere below as a "basic apparatus".

In the aforementioned particular case of using a pulse-dial telephone set as a registration-playing terminal, the processor (3) is a mini-automatic telephone exchange designed for, as a minimum, N telephone numbers... ..a voice generator under a preset program. In this case, N telephone numbers of the processor (3) are used to receive signals about wagers placed. When answering a subscriber's call by one of these N numbers, the processor (3) generates a signal containing a speaker's telephone number as identification data and an... ..a game set element which corresponds to a dialed telephone number of the input/output processor (3).

A quantity of numbers in a telephone exchange for a game among N-power... ...for rendering additional services to the players. For example, (N+1) telephone number of the processor (3) may be used for receiving inquiries to generate a wager from the players, (N+2) telephone number of the processor (3) may be used for purchasing a credit to be entitled to wager, (N+3) telephone number of the processor (3) may be used for informing about results of the players in completed round for a certain period of time, (N+4) telephone number of the processor (3) may be used for returning wagers from unfinished playing rounds, etc. So, when answering a subscriber's call by the telephone number of an additional service, the processor (3) generates a signal containing a speaker's telephone number and a special code of this service.

An output of the processor (3) containing a speaker's telephone number as identification data and an information code of... ...an additional service code which corresponds to a dialed telephone number of the input/output processor (3), is recognized and identified in the recognition and identification unit (4). In the course of signal identification, a speaker's telephone number is associated with a database address in the long-term memory unit (14) under which a personal account-register is stored, if any, or a speaker receives... ...to dial a (N+2) telephone number in order to open an account-register and purchase a minimum credit for the right to wager. After the signal containing a wager value has been identified, this signal is processed according to the above-described scheme, wherein data generated by the wager registration confirmation unit (12) and transferred to the processor (3) are converted by the voice generator and distributed to a speaker in response to... ...wager is placed by a wager generator (13). Upon identification of a special playing credit purchase code, the payment registration unit (15) corrects accordingly a speaker's account-register balance in the long-term memory unit (14) and produces to the processor (3) data about the corrected account balance which, after conversion by the voice generator in the processor (3), is distributed to the subscriber in a voice form in response to his or... ...Upon identification of a special code of a playing round outcome review, the outcome review unit (16) analyzes a speaker's account-register and produces to the processor (3) data about the outcome of a subscriber's wager drawing for a certain period of time which, after conversion by the voice generator in the processor (3), is distributed to the subscriber in a voice form in response to his or her call, etc. Essentially, payments for the participation in the game using the described modification of the apparatus are effected according to one of the known procedures for paying telephone services.

To implement... ...request, with information about a current wager distribution among the game set elements, the basic apparatus is supplemented with a wager drawing display unit (18) coupled between an additional output of the controller (7) and an additional input of the processor (3). To obtain information about a current quantitative wager distribution among the game set elements, that is about a current wager drawing state, the player sends from his or her terminal to the processor (3) a signal carrying, along with identification data of the player, a data inquiry flag about the current drawing state. When the recognition and identification unit (4) detects in the signal received from the processor (3) an identified player's inquiry flag to information about the current drawing state, the recognition and identification unit (4) inquires on the wager generator (13) for a wager and supplements the signal received from the

processor (3) with its value, thereupon transfers the signal so converted to the wager payment unit (5) in order to draw on an account of the identified player in the amount... ...Thereupon, a signal containing information about a wager placed without participation of the player is processed by other functional units; in the course of processing the signal with an inquiry flag to present information the controller (7) produces to the wager drawing display unit (18) data about quantitative wager distribution among game set elements and data about revealed special points of this distribution which are converted by the unit (18) into a message format about a wager drawing current state and transfers this message to the processor (3) for communication to the player who has sent the inquiry.

Return of wagers may be only conducted in the games with a postponed outcome of the drawing process from unfinished playing rounds.

To implement another particular case of a wagering game method in... ..the present invention which enables the players to withdraw their wagers from a wager drawing process in games with a postponed outcome of the drawing, the basic apparatus is supplemented with a wager returning unit (19) coupled between the recognition and identification unit (4) and the input/output processor (3) and additionally interconnected with the controller (7) and the long-term memory unit (14) (these two connections are shown in FIG. 1 with a dotted line). To return... ..wagers from an unfinished playing round, the player sends from his or her registration-playing terminal to the processor (3) a signal carrying, along with identification data of the player, a request flag for wager return. In the recognition and identification unit (4), such a signal is identified and recognized as a signal of identified player for returning his or her wagers and transferred to the wager returning unit (19) which first of all checks, with the help of the long-term memory unit (14), the availability in an account-register of this player of data about registered wagers... ..the wager return signal is ignored, whereas with a positive check result the wager returning unit (19) initiates data exchange with the controller (7) providing the latter with a database address... ..for wager return. This address is transferred by the controller (7) via a playing-logic unit (8) to a recording unit (9) for entering in the protocol of the current playing round. The controller (7) interrupts reception of signals with wagers from the wager registration unit (6) and changes over to interaction with a wager returning unit (19). The wager returning unit (19) receives from the long-term memory unit (14) and transfers to the controller (7) information codes of wagers placed by a given... ..code of the regular wager requested for return, the controller (7) transfers this code for processing by a wager distribution processor (11). The wager distribution processor (11) processes a wager code received from the controller (7) in, the usual fashion, except for that... ..code of a corresponding game set element, a number of correlated wagers accumulated in the processor (11) is not increased by 1, but rather decreases by 1. As a result of processing the regular wager with a return flag, the processor (11) produces, as usual, a signal-flag F indicative of the completion of wager drawing... ..the current round, said signal being translated by the controller (7) to the wager returning unit (19).

Having received from the processor (11) a signal $F=0$, after processing of the regular wager code with a return flag, the controller (7) transfers this wager code via the playing-logic unit (8) to the recording unit (9) in order to enter in the current round protocol a mark indicative of the... ..was entered in the protocol earlier. Concurrently, the flag $F=0$ for the wager returning unit (19) is a confirmation of the regular wager return with a return flag and indication... ..a return flag. An account-register of the player located in the long-term memory unit (14) is added with a wager return mark along with a correction of an account- register balance. Having received from the processor (11) a signal $F=1$, after processing of the regular wager code with a return flag, the controller (7) carries the current... ..to completion in the usual fashion, while unloading the buffer storage via the playing-logic unit (8) to the recording unit (9); wagers returned are disregarded in the calculation of a completed round budget and a... ..allocation. Concurrently, signal $F=1$ completes interaction of the controller (7) with the wager returning unit (19), so that not a single wager from the completed round comes back. On completion of a wager return session, the unit (19) via the processor (3) produces for a player who requested a wager return, a message about which wagers... ..method in accordance with the present invention, the controller (7) interacts with the wager distribution processor (11) which is assembled in accordance with a function and logic circuit represented in FIG. 2.

The controller (7) sends to the processor (11) a wager information code as a set of binary number bits which enters an... ..22) at the F output of which on-bit emerges, if, and only if the processor (11) processes a wager placed on the last "unoccupied" game set element. An output of the gate... ..to reset inputs of all flip-flops (21) and also as an output of the processor (11) - to the controller (7). Thus, signal $F=0$ does not change the states of... ..and serves for the controller (7) as evidence of "failure to win" of a wager processed, so that the controller (7) must proceed with the current playing round, whereas signal F... ..re-sets all flip-flops (21) to 0, thus preparing the controller (7) for wager processing in the next round, and gives instructions to the controller (7) about the current round... ..depending upon the rules of a prize fund allocation to be used, the playing-logic unit (8) and the recording unit (9) need data about a precise quantitative wager distribution among the game set elements, then in order to ensure this, the wager distribution processor (11) may be assembled in accordance with a function and logic circuit represented

in FIG... flops, use is made of N binary counters (23-1), ..., (23-N) with null-comparison units (24-1), ..., (24-N), wherein each 1-bit output of the decoder (20) is an... element. A reading of each counter (23) is supplied to an input of the null-comparison unit (24), so that if the input coincides with 0, an output of the null-comparison unit (24) is set to 1, if it does not coincide - to 0. Outputs of the comparison units (24) are supplied to a "logical AND" gate (25) with N inverse inputs, whose output F is the output of the processor (11) to the controller (7). In so doing, $F=1$ if, and only if the processor (11) processes a code of the wager placed on the last "unoccupied" game set element. Besides, signal... playing round. Signal $F=1$ re-sets all counters (23) to 0, thus preparing the processor (11) for wager processing in the next round, and serves to the controller (7) an indication about the current... 7) over a data bus (26) under control of the controller (7).

A wager distribution processor in accordance with an apparatus realizing a "Force of Minimum" game method may be assembled according to a circuit depicted... set element, and readings of the counter are additionally compared with 1 in a comparison unit (27). When an input binary number coincides with 1, the bit of a comparison unit (27) sets to 1, when it does not coincide - to 0. Outputs of comparison units (27) are supplied to a N-input "exclusive OR" gate (28), whose output F_1 takes... 1 when, and only when exactly one input equals 1. Besides, outputs of all comparison units (27) are supplied to N inputs of an encoder-address former (30). Signals of null-comparison units (24) are processed, like in the scheme depicted in FIG. 3, by a "logical AND" gate (25) with... $F=1$ re-sets the counters (23) to an initial zero state, thus preparing the processor for wager processing in the next round. Thus, emergence of signal $F=1$ at the output of the wager distribution processor (11) is a flag for the controller (7) about the current round completion, wherein the... signal not to increase, but to decrease an accumulated sum by 1.

A wager distribution processor in accordance with an apparatus realizing a "Force of Minimax" game method may be assembled according to a circuit depicted... also corresponds to each game set element, however an output of every such counter is processed by three comparison units: (24), (31), and (32). The unit (24) carried out comparison with 0, the unit (31) - comparison with an output of a minimum-counter (33), the unit (32) - comparison with an output of a maximum-counter (34). Initialization of all counters is carried out by the processor's output $F=1$ upon completion of the regular round; in so doing, all N... is set equal to 3. Thus, in the beginning of each round the minimum-comparison unit (31) compares readings of "element" counters (23) with number 1 (a value of the first nonzero global minimum), the maximum-comparison unit (32) - with number 3 (a value of the first global maximum given a nonzero global minimum).. Outputs of a null-comparison unit (24), as before, are processed by a "logical AND" gate (25) with N inverse inputs, whose output takes a value... distribution, and a value $F_0=0$ if they are present. Outputs of the minimum-comparison unit (31) are supplied to corresponding inputs of a first encoder (30-1) and, besides, are processed by a first N-input "exclusive OR" gate (28-1). Outputs of the maximum-comparison unit (32) are supplied to corresponding inputs of a second encoder (30-2) and, besides, are processed by a second N-input "exclusive OR" gate (28-2). As a result, the output... 1 if, and only if exactly one coincidence has been fixed in the minimum-comparison

units (31), and the output F_2 of the second gate (28-2) takes a value of 1 if, and only if exactly one coincidence has been fixed in the maximum-comparison units (32). Signals F_0 and F_1 are supplied to a first two-input "logical AND" gate in the wager distribution.

As a result, the wager distribution processor (11) transfers to the controller (7) as its output data a binary number at the... controller (7) is able to read data about the current wager distribution out of the processor (11), via the data bus (26).

The function of a wager withdrawal from the current... by the counters will not increase, but decrease by 1.

Modifications of the wager distribution processor (11) as described above may be implemented differently: in a form of custom and semi-custom integrated circuits, specialized computer cards made of conventional components,

on the basis of specially dedicated main memory arrays of the personal computer under control of an application program, etc.

Industrial applicability

The following examples deal with some... ..are called, for short, iterative-analytical games, thus reflecting the principle of a wager drawing process constituting the basis of these games.

Example 1. Games In the Slot Machine Centers

Iterative-analytical games may be conducted using playing terminals in casinos and specialized slot machine centers. Traditionally, slot machines embody various "face-to-face... ..When using the claimed method, a slot machine offers the services of a registration-playing terminal with respect to the apparatus described above.

Some apparatuses embodying different iterative-analytical games ("Force of Zero", "Force of Minimum... ..Minimax") among game sets of different power N and some slot machines and electronic cash terminals of the players' registration and wager payment may be integrated through a local computer network under control of a computer playing server into a playing system offering the players ample scope of choice. Moreover, in such a system games identical by an iterative-analytical process and power N of a game set may differ by a price of wagers to be placed. In so doing, a playing server in respect of each apparatus will act as the input/output processor (3), whereas a server database - as the long-term memory unit (14) (see FIG.1).

Before entering into the game, the player notifies the operator of the cash terminal of required personal data for registration, pays a necessary sum of the playing credit and receives from the playing terminal a registration card with a machine-readable data medium containing an individual code which is...her with information about the current state of wager drawing which are sent by a computer server for processing in a corresponding apparatus embodying a game type chosen.

For convenience, slot machines may be equipped with printers to... ..players of wagers placed by them, whereas a local network may additionally comprise reference-information terminals through which everybody who wishes so may obtain information about completed playing rounds. Any player having produced his or her registration card may receive, through cash terminal operators, a money equivalent of the cost of a playing credit balance from his or... ..the playing credit which will be entered to his or her account. Moreover, using cash terminal operators, a player who is not able to wait for the completion of rounds with... ..or her wagers from unfinished playing rounds. Upon receipt of such a request, a playing server blocks a further acceptance of wagers and requests associated with an individual card code of... ..a local area network of a "slot machine centers" may contain only 1-2 playing terminals; however, it may be added with a communication input/output processor supporting interaction of the playing server with remote playing terminals through a switched communication line (direct-connection telephone service) and/or batches (TCP/IP - Internet... ..center" can participate in the game from any other place with the available telephone and computer with modem or, at least, a touch-tone telephone. If the player has access to Internet, the communication input/output processor provides him or her with a whole range of services of a "slot machine center... ..credit cards; in the absence of access to Internet and with the availability of a computer with modem, the same services may be rendered through a direct telephone connection in the... ..hotels in which rooms are equipped with a cable system for broadcasting TV programs requiring payment or intrahotel cable channel for providing information services in the interactive mode.

A guest, using... ..TV screen menu which in this particular case offers the services of a registration-playing terminal, carries on a dialog with a playing system such as if it were in the... ..won are materialized by the hotel in the money or other form (souvenirs, free-of-charge services, free nourishment, living-conditions enhancement, etc.).

When discharging from the hotel, a guest may... ..an authorized user of its services with a personal network address, registration number and adjusted payment system, so that the opening of a personal paying account and correction of its balance... ..case, instead of the play for money, a game organizer may assign wager prices in units for

measuring consumer-payable information resources to be supplied by the network (for example, in the minutes of sport or entertainment program broadcasting time, optionally), and settle accounts with players just in these units using part of resources released as a result of the game conduct (but already paid...sales of numbered lottery coupons to the population (like "Lotto-Million", see above) with electronic terminals for wager registration, may be used, without substantial alterations, for the conduct of iterative-analytical... ..numeric value of this wager), as well as to hardware-software of the central playing server which must be adapted to a wager registration mode with timing separation in order to... ..embody iterative-analytical games. When providing a necessary speed of response of the central playing server and sufficient channel capacity, an electronic terminal of the point of sales of lottery coupons prints on a player's coupon during... ..a lottery information service. In case of a "Force of Zero" instantaneous lottery, the electronic terminal prints an outcome on a player's coupon right after a wager has been read from the coupon and registered by the central server in a corresponding apparatus in accordance with the present invention which embodies a selected game.

To simplify operations concerning... ..lottery games may be conducted according to a multilevel hierarchical scheme of connecting the playing servers, where each level corresponds to a certain range of power N of game sets, so that, for example, $N \leq 100$ games are processed by servers covering territories with population up to 20-30 thousand people, $N \leq 1000$ games - by servers covering territories with population up to 100-150 thousand people, $N \leq 10,000$ games - by servers covering territories with population up to 1.5-3.0 million people, etc. Provision of electron wager registration terminals in the point of sales of lottery coupons with the screens to display the current... ..value interactive mode of conducting iterative-analytical games, thereby making them more attractive to mass consumers in comparison with the existing national lotteries.

Example 4. Telephone Lotteries and TV-Broadcast Games... ..in many countries for the conduction of telephone lotteries using touch-tone telephone sets and computer servers carrying out registration of players and wager acceptance by means of a successive generation of... ..answers to these messages in a form of digital codes to be generated by a server when identifying signals sent by a speaker from a touch-tone telephone, may be used with success for the conduction of iterative-analytical games.

Technical changes affect only a playing server and comprises adaptation of voice generation programs to the specific nature of iterative-analytical games... ..number in the turn and the number of a playing round, etc.), introduction in the server hardware-software of a wager registration mode with timing separation in order to form turns and its supplementation with iterative-analytical processors, as well as adaptation of the used database of those participating in the lottery.

In... ..settlements between the game participants and authorities may be the assigning of wager prices in units for measuring services rendered by the telephone network (for example, in the minutes of a... ..the duration of local, trunk or international calls for which a player is exempted from payment), and a telephone network is able to use its own part of revenue from the... ..information about the current playing round state is ensured to those players having a home computer with modem which may interact with a playing server through Internet or in the direct connection mode. Those players lacking a computer with modem but being subscribers of paging networks, may order transfer of information about a... ..to their pagers. Finally, in agreement with TV broadcasting stations, the course of wager drawing processes may be the subject of regular broadcast on the air to TV receiving antennas; during such broadcast, when any wager comes from the players for registration, the server may withdraw from ...such telecasts of additional orders for advertising.

With a special arrangement of the input/output processor (3) of apparatuses for realizing iterative-analytical games in accordance with the present invention, telephone... ..a game wager. The players send the completed forms with data about wagers to the server of the game authority, where this data is registered in real time and processed in a corresponding apparatus in accordance with the present invention. In so doing, the

server provides the player with a wager registration message containing a regular advertisement which will be...
...network and from above examples it is clear that the adjustment of electronic system of payments through credit cards and possibility of rapid exchange of both text and graphical information makes... ...and companies, it is possible to consider the possibility to conduct iterative-analytical games using servers of those entities providing the Internet services (Internet-providers) with payment of wagers in units for measuring the cost of providers' services (for example, in minutes of the Internet access... ...to offer its services to certain groups of the population and institutions on free-of-charge or preferential terms. This circumstance may promote the drawing of further participants in the game... ...games by

Internet-providers into a socially useful measure.

Example 7. Games In Discount and Payment Systems

In the 90s, consumers' markets of many countries were flooded with discount systems designed for the drawing regular customers in large provides of goods and services at the expense of providing these customers with flexible systems of individual rebates (discounts) for goods and services and of other different privileges, where an count of each customer is supplemented with a certain amount of bonus points for each payment received from him or her (or a sum of payments for a certain period of time) followed by realization of accumulated points through the aforementioned... ...become possible thanks to the introduction by providers of goods and services of hierarchical corporate computer networks covering each cash register through which the sale of a good or service is carried out, and containing terminals for the registration of new customers in all points of the provider - customer interaction. In the course of registration in the discount system, each customer receives a personal identification number which is applied onto a personal discount card to be handed out to a customer and which is associated in a one-to-one manner with a system file to be created for a customer's personal data and his or her personal account-register for chronological account of all payments, bonus points added and also discounts and other privileges used by the customer (see, for example, Internet resources <http://www.gb.be> of the Belgian supermarket network GB... ...gas station network Petro-Canada, <http://www.transaero.ru> of the Russian airline Transaero).

Internet-servers of discount systems not only provide a detailed information about discounts and privileges but also start providing the customer with access via a PIN-code to his or her personal account-register of bonus...
...introduction consists in the fixing of wager prices in bonus points. Indeed, if the Internet-server of the discount system has an application with iterative-analytical games installed, then any participant... ...seats (see, Example 2, as well as Swissair Gazette, June'98, p.103).

Numerous bank payment systems by plastic debit and credit cards and also by cards with a built-in... ...Proton), are close to discount systems in respect of the organization and technological infrastructure used. Payments by debit and credit cards are widely spread in medium and large retailers and services, payments with smart-cards are used to pay for everyday petty expenses: urban transport travel, parking time, calls in the public telephone, buying newspapers, etc. Holders of such cards have been identified in the system with the numbers... ...in smart-cards) is supported by networks of automatic teller machines each representing a personal computer with a graphic controller. For a special purpose to load smart-cards, there have been created networks of additional electronic terminals with alphanumeric displays which are built in public touch-tone telephones booths or may be... ...money without going out of doors.

Interaction-analytical games as one of applications of electronic payment systems through debit, credit and smart cards, due to their playing interest and sporting competitive data about the game set elements. The player purchases a distributor's playing coupon, marks in its information block a game set element on... ...of 32 wagers.

When embodying this wagering game method, use may be made of the computer technology, in particular special computer programs and systems of reading data out of paper carriers, thus enabling a rather quick reading of identification data and, marks about the placed wagers from coupons and processing of these data and marks by an algorithm which embodies a wagering game method in... ...a game organizer, in the order of arrival priority, in an

iterative-analytical coupon drawing process, while registering for each coupon a serial wager number in the current playing round and... ..N. To form the turn of coupons-wagers and run an iterative-analytical coupon drawing process, different methods and means may be used, ranging from a manual sorting and arrangement of... ..table under control of a "counting committee" to the application of playing systems equipped with computerized terminals of data reading, recognition and registration, with the implementation of drawing processes through an application package or special processors similar to those proposed by the present invention. In case of using computer systems, it is advisable to represent outcomes and processes of wager drawing on screens or electronic displays during breaks and on completion of mass... ..enhance the interest potential players in the game.

So, the performance of an iterative-analytical process over a waiting line of wagers ensures a full dependence of the game outcomes upon... ..retaining profitability for a game organizer, to create playing applications so attractive for a mass consumer ranging from instantaneous "paper" or electronic lotteries and slot machines to national interactive telephone lotteries...

Claims: ...by means of generation of a set of N non-repeating information codes in a computer memory, propagation, among the players through communication lines, of signals carrying information about the game... ..the wagers, identification and registration of signals received through feedback lines, forming of a wager payment data, a wager drawing within playing rounds, characterized in that the signals carrying wager information... ..is completed, and the wager drawing is carried out by means of an iterative-analytical process of forming a quantitative wager distribution among the game set elements, said process is kept hidden from the players until the playing round is completed, and within every iteration of the said process a regular signal of a registered signal sequence is correlated with the information code of... ..said conditions are observed, and in the presence of registered wager information carrying signals not processed by the iterative-analytical wager drawing process before the completion of the current playing round the said signals are processed by the iterative-analytical wager drawing process within one of the next rounds.

2. A method of claim 1, characterized in that within every iteration of a wager drawing process, information codes are revealed with which no signal has been correlated within the current playing round, and the iterative-analytical wager drawing process is completed in the processing of a signal containing information about a wager on the only game set element with whose information code no signal has been correlated by the iterative-analytical process within the current playing round before the processing of this signal.

3. A method of claim 1, characterized in that within every iteration of a wager drawing process, starting with the N iteration, information codes are revealed with which no signal has been correlated within the current playing round, and the iterative-analytical wager drawing process is completed in the processing of a signal containing information about a wager on the only game set element with whose information code no signal has been correlated by the iterative-analytical process within the current playing round before the processing of this signal.

4. A method of claim 1, characterized in that within every iteration of a wager drawing process, information codes are revealed with which only one signal has been correlated within the current playing round, and the iterative-analytical wager drawing process is completed in the presence of only one said information code and in the absence of information codes with which no signal has been correlated by the iterative-analytical process within the current playing round before the processing of this signal.

5. A method of claim 1, characterized in that within every iteration of a wager drawing process, starting with the (2N-1) iteration, information codes are revealed with which only one signal has been correlated within the current playing round, and the iterative-analytical wager drawing process is completed in the presence of only one said information code and in the absence of information codes with which no signal has been correlated by the iterative-analytical process within the current playing round before the processing of this signal.

6. A method of claim 1, characterized in that with a game set formed by $N \geq 2$ information elements, within every iteration of a wager drawing process, information codes are revealed with which the minimum and maximum

number of signals have been correlated within the current playing round, and the iterative-analytical wager drawing process is completed in the presence of only one information code with which the minimum number... ..set formed by $N+2$ information elements, within every iteration of a wager drawing process, starting with the $2N$ iteration, information codes are revealed with which the minimum and the... ..signals have been correlated within the current playing round, and the iterative-analytical wager drawing

process is completed in the presence of only one information code with which the minimum number...is completed, and the wager drawing is carried out by means of an iterative-analytical process of forming a quantitative wager distribution among the game set elements, said process is kept hidden from the players until the playing round is completed, and within every iteration of the said process a regular coupon of a registered coupon sequence is correlated with the game set information... ..are observed, and in the presence of registered coupons carrying wager marks which are not processed by the iterative-analytical wager drawing process before the completion of the current playing round, these coupons are processed by the iterative-analytical wager drawing process within one of the next rounds.

9. A method of claim 8, characterized in that within every iteration of a wager drawing process, the game set elements are revealed with which no coupon has been correlated within the current playing round, and the wager drawing is completed with an iterative-analytical processing of a coupon containing a wager mark corresponding to the only game set element with which no coupon has been correlated within the current playing round before the processing of this coupon.

10. A method of claim 8, characterized in that within every iteration of a wager drawing process, starting with the N iteration, the game set elements are revealed with which no coupon... ..within the current playing round, and the wager drawing is completed with an iterative-analytical processing of a coupon containing a wager mark corresponding to the only game set element with which no coupon has been correlated within the current playing round before the processing of this coupon.

11. A method of claim 8, characterized in that within every iteration of a wager drawing process, the game set elements are revealed with which only one coupon has been correlated within... ..A method of claim 8, characterized in that within every iteration of a wager drawing process, starting with the $(2N-1)$ iteration, the game set elements are revealed with which only... ..set formed by $N+2$ information elements, within every iteration of a wager drawing process the game set elements are revealed with which the minimum and the maximum number of coupons have been correlated within the current playing round, and the wager drawing process is completed in the presence of only one element with which the minimum number of... ..set formed by $N+2$ information elements, within every iteration of a wager drawing process, starting with the $2N$ iteration, the game set elements are revealed with which the minimum... ..number of coupons have been correlated within the current playing round, and the wager drawing process is completed in the presence of only one element with which the minimum number of... ..player in exchange for a wager which is placed without his or her participation and processed by a wager drawing process out of turn.

16. A method of claims 1-14, characterized in that, on request... ..wager information which were received from the said player are withdrawn from a wager drawing process in the order opposite to that of their registration.

17. A wagering game apparatus to carry out a method as claimed in claim 1, comprising a game set forming unit (1) connected via data dissemination unit (2) to one of inputs of a processor (3) connected with its information output to a recognition and identification unit (4), a wager payment unit (5), a wager registration unit (6), a controller (7), a playing-logic unit (8), and a recording unit (9) which are connected in series, a playing-round counter (10) connected to the second input of the wager registration unit (6) and to the second output of the controller (7) connected with its second input to the output of the game set forming unit (1), a long-term memory unit (14) interconnected with the recognition and identification unit (4) and the wager payment unit (5), a timer

(17) connected to the controller (7), the recognition and identification unit (4), the wager payment unit (5), and the recording unit (9), characterized in that it further comprises a wager distribution processor (11) interconnected with the controller (7), a wager registration confirmation unit (12) connected to the input of the processor (3) and the second output of the wager registration unit (6), a payment registration unit (15) and an outcome review unit (16) which are interconnected with the long-term memory unit (14) and the processor (3) and also connected to corresponding outputs of the recognition and identification unit (4), the outputs of the recording unit (9) and the wager registration confirmation unit (12) being connected to corresponding inputs of the long-term memory unit (14).

18. An apparatus of claim 17, characterized in that it comprises a wager generator (13) interconnected with the recognition and identification unit (4) and also connected to the output of the game set forming unit (1).

19. An apparatus of claims 17-18, characterized in that it comprises a wager drawing display unit (18) coupled between the controller (7) and the input/output processor (3).

20. An apparatus of claims 17-18, characterized in that it comprises a wager returning unit (19) interconnected with the controller (7) and the long-term memory unit (14) and also connected to an output of the recognition and identification unit (4) and an input of the input/output processor (3).

21. An apparatus of claims 17-20, characterized in that the wager distribution processor (11) comprises a decoder (20) connected with its outputs to driving inputs of flip-flops... ..22) connected with its output to reset inputs of the flip-flops (21).

22. An apparatus of claims 17-20, characterized in that the wager distribution processor (11) comprises a decoder (20) connected with its outputs to inputs of counters (23) whose outputs are connected, via comparison units (24), to inverse inputs of a "logical AND" gate (25) connected with its output to reset inputs of the counters (23).

23. An apparatus of claims 17-20, characterized in that the wager distribution processor (11) comprises a decoder (20), each of N 1-bit outputs of the said decoder being coupled to a stage of a counter (23) and of a null-comparison unit (24) and a 1-comparison unit (27) which are connected in parallel to the said counter (23), a "logical AND" gate (25) with N inverse inputs each coupled to the output of the corresponding null-comparison unit (24), an "exclusive OR" gate (28) with N inputs each coupled to the output of the corresponding 1-comparison unit (27), a "logical AND" gate (29) with two inputs connected to outputs of the gates... ..encoder (30) with N inputs each coupled to the output of the corresponding 1-comparison unit (27), the said gate (29) being connected with its output to reset inputs of the counters (23) and to a control input of the encoder (30).

24. An apparatus of claims 17-20, characterized in that the wager distribution processor (11) comprises a decoder (20), each of N 1-bit outputs of the said decoder being coupled to a stage of a counter (23) and of a null-comparison unit (24), a minimum-comparison unit (31), and a maximum-comparison unit (32) which are connected in parallel to the said counter (23), a "logical AND" gate (25) with N inverse inputs each coupled to the output of the corresponding null-comparison unit (24), a first "exclusive OR" gate (28-1) with N inputs each coupled to the output of the corresponding minimum-comparison unit (31), a first "logical AND" gate (29-1) with two inputs connected to outputs of... ..28-2) with N inputs each coupled to the output of the corresponding maximum-comparison unit (32), a second "logical AND" gate (29-2) with two inputs connected to outputs of... ..30-1) with N inputs each coupled to the output of the corresponding minimum-comparison unit (31), a second encoder (30-2) with N inputs each coupled to the output of the corresponding maximum-comparison unit (32), a minimum-counter (33) coupled to the output of the gate (29-1), the... ..counter (33) being connected with its output to the input of each of minimum-comparison units (31), and a maximum-counter (34) coupled to the output of the gate (28-2)... ..counter (34) being connected with its output to the input of each of maximum-comparison units (32), the said gate (29-2) being connected with its output to reset inputs of... ..with its output to a control input of the second encoder (30-2).

25. An apparatus of claims 17-24, characterized in that the input/output processor (3) includes a telephone exchange for at least N telephone numbers with an automatic speaker's telephone number determinant and a controlled voice generator.

26. An apparatus of claims 17-24, characterized in that the input/output processor (3) includes a computer network server and a unit for contacting clients of the said network.

6/K/2 (Item 2 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Country	Number	Kind	Date		
Type		Pub. Date		Kind	Text
Available Text		Language		Update	Word Count
Total Word Count (Document A)					
Total Word Count (Document B)					
Total Word Count (All Documents)					

Specification: ...by reference. BACKGROUND OF THE INVENTION

The present invention relates to an electronic cashing card settlement system for electronic money such as an IC card, a prepaid card or the like. In particular, the present invention relates to an electronic cashing card settlement system in which two kinds of money processing areas are prepared depending on input or non-input of a password number within a... ..illustrated in FIGS. 9 - 11 has been proposed as an electronic cashing card for transaction settlement. FIG. 9A is a schematic block diagram of an IC card used as an electronic... ..an electronic cashing card 100 of the related art comprises a non-personal authentication money processing memory 101 for executing settlement of money without request for a password number for personal identification. A personal authentication money processing memory 102 is for storing data for personal authentication money processing to execute settlement of money responding to a request for such a password. Data write/read controlling means... ..controlling data writing to or reading from the non-personal authentication and personal authentication money processing memories 101, 102. Controlled arithmetic operation means 104 is for executing settlement of money for non-personal and personal authentication money processing and various kinds of controlled arithmetic operations. Input/output means 107 is for executing data input/output between the controlled arithmetic operation means 104 and a read/write (R/W) unit or ATM (Automatic Teller Machine) of a bank (not illustrated), to which this IC card 100 is inserted.

FIG. 10A illustrates the non-personal authentication money processing memory 101 explained above. Memory 101 comprises a regional code of a management organization, a... ..storing data such as personal information or the like. Further, a non-personal authentication money processing area 101b allows entry of deposit, disbursement, and balance of money as a history in regard to non-personal authentication money processing. Area 101b also allows update of such data.

Meanwhile, the personal authentication money processing memory 102 is illustrated in FIG. 10B. Memory 102 comprises, like the non-personal authentication money processing memory 101, an ordinary data area 102a (corresponding to 101a in FIG. 10A) and a personal authentication money processing area 102b (corresponding to

101b in FIG. 10A) in common. In addition to this structure... ..a password number in the personal information of the ordinary data area 102a.

Next, a settlement operation of an electronic cashing card of the related art, based on the structure explained above, will be explained with reference to FIG. 11. FIG. 11 illustrates a settlement operation flowchart in the non-personal authentication and personal authentication money processing of an electronic cashing card as illustrated in FIG. 9A.

In FIG. 11, a consumer who owns an electronic cashing card shops at a store. An amount of payment for the items selected for purchase is determined as an amount of sales from a terminal equipment (not illustrated) of a Point Of Sales System (POS) (step SP 101). The consumer determines which of the non-personal or personal authentication money processing should be used for this payment, depending on input or non-input of the password number (step SP 102).

When It is determined to execute the settlement without input of the password number, settlement is executed by the non-personal authentication money processing. A non-authentication balance 101c in the non-personal authentication money processing area 101b is read via the write/read control means 103 (step SP 103). This... ..balance 101c is determined to be larger than or equal to the amount of sales, payment by the non-personal authentication money processing area 101b is executed (step SP 105).

After the settlement by the non-personal authentication money processing is executed, the non-authentication balance 101c is updated (step SP 106). A new amount... ..is determined to be smaller than the amount of sales in step SP 104, a process for disabling settlement of the transaction by the non-personal authentication money processing is executed (step SP 108). This disablement process is executed, for example, by display or announcement of a money shortage amount.

On the other hand, when it is judged that settlement is executed by input of the password in step SP 102, settlement is executed by the personal authentication money processing. In this personal authentication money processing, the password number data in the ordinary data area 102a is read via the write... ..is determined whether the password number data matches with the password number input by the consumer (step SP 110).

When a match of the password number data is found at step SP 110, payment by the personal authentication money processing area 102b is executed (step SP 111). After the settlement by the personal authentication money processing is executed, the authentication balance 102c is updated (step SP 112) and a new amount... ..a failure to match the password number data is found at step SP 110, a process for disabling settlement of the transaction by the personal authentication money processing is executed (step SP 114).

Moreover, an example of another electronic cashing card of the... ..designed as a type of prepaid card. Card 200 comprises a non-personal authentication money processing memory 201, for writing the amount of money set by the pre-payment in regard to the non-personal authentication money processing. This non-personal processing is for executing settlement of money without request for a password number as personal authentication. Memory 201 also is for writing the balance after the amount of settlement is subtracted from such preset amount.

Moreover, a personal authentication money processing memory 202 is for writing the amount of money set by a pre-payment in regard to the personal authentication money processing. This personal processing is for executing settlement of money with a request for the password number. The balance after the amount of settlement is subtracted from such preset amount.

These non-personal authentication and personal authentication money processing memories 201, 202 are formed in a structure such that a magnetic recording tape is adhered to or buried in the side surface of the card. After the settlement, data writing to the non-personal authentication and personal authentication money processing memories 201, 202 is executed in the same manner as the IC card type electronic... ..card of the related art is structured as explained above, the non-personal authentication money processing and personal-authentication money processing

are executed individually and independently of each other. This individual execution applies in both the... ..type and the case of the prepaid card type.

The problem therefore results that a settlement cannot be made when the amount of money has exceeded the amount preset by the non-personal authentication money processing. Namely, an amount of money exceeding the preset limit amount cannot be settled, because the allowable limit amount is generally set from the point of view of transaction security for the non-personal authentication money processing, wherein the password number is not requested. This is the case even when settlement can be executed by cash or deposit to the account of the personal authentication money processing in which the password number is requested.

In this case, it is also considered to conduct a deposit process by shifting money for the personal authentication money processing to money for the non-personal authentication money processing. However, such a deposit process can be executed only with an ATM of a bank or by specially designed R/W equipment. Therefore, the problem arises that settlement at the store location, where shopping is done, is disabled. The range of applications for the card is thereby limited, and... ..the amount of sales exceeds the preset limit amount for the non-personal authentication money processing, but it is possible to make a deposit by shifting money from the personal authentication money processing. In this case a useful purpose is served by previously requesting the password number in the personal authentication money processing. However, the option of such a deposit, even through an authentication process where a password number is requested, has resulted in the problem that transaction security cannot... ..the state of the art in the introduction refers to electronic cashing cards and the settlement of monetary transactions, and references to aspects of the invention and the description of exemplary... ..of the present invention given below, and claims, also refer to cashing cards and the settlement of monetary transactions, it should be understood that the present invention finds application in other... ..balance information, other than money, are amounts of time or amounts of other tokens or units which may be used in transactions involving those quantities. Units of time may, for example, be involved in transactions involving the use or hire of... ..use when travelling on a public transport network, and one example of another kind of unit which may be involved is telephone call units. There is no restriction on the quantities which may be represented by balance information. The... ..merely to refer to currency or money as such, but also to any tokens or units which can be "spent" or used by a card user.

According to the present invention there is provided a system for transaction settlement with an electronic cashing card having a non-authentication processing memory and an authentication processing memory, said system comprising:

means for updating an authentication balance stored in a balance area of the authentication

processing memory and a non-authentication balance stored in a balance area of the non-authentication processing memory, said means updating the authentication balance to a balance amount after settlement when a transaction is settled by an authentication process having a requirement for a personal authentication to be matched, said means updating the non... ..an amount less than or equal to the stored authentication balance when the transaction is settled by the authentication process; and

means for comparing the non-authentication balance and the authentication balance and determining that an illegal process has been performed with the card when the non-authentication balance is larger than the authentication balance.

According to the present invention there is also provided a method of transaction settlement with an electronic cashing card having a non-authentication processing memory and an authentication processing memory, the method comprising:

updating an authentication balance stored in the authentication processing memory and a non-authentication balance stored in the non-authentication processing memory, the authentication balance being updated to a balance amount

after settlement when a transaction is settled by an authentication process having a requirement for a personal authentication to be matched, the non-authentication balance being... ..an amount less than or equal to the stored authentication balance when the transaction is settled by the authentication process; and

determining that an illegal process has been performed with the card when a comparison of the non-authentication balance and... ..than the authentication balance.

According to the present invention there is provided a program for settlement of transactions with an electronic cashing card having a non-authentication processing memory and an authentication processing memory, said program comprising procedures for:

updating an authentication balance stored in the authentication processing memory and a non-authentication balance stored in the non-authentication processing memory, the authentication balance being updated to a balance amount after settlement when a transaction is settled by an authentication process having a requirement for a personal authentication to be matched, the non-authentication balance being... ..an amount less than or equal to the stored authentication balance when the transaction is settled by the authentication process; and

determining that in illegal process has been performed with the card when a comparison of the non-authentication balance and... ..is larger than the authentication balance.

According to the present invention there is provided a computer readable medium encoded with a program for settlement of transactions with an electronic cashing card having a non-authentication processing memory and an authentication processing memory, said program comprising procedures for:

updating an authentication balance stored in the authentication processing memory and a non-authentication balance stored in the non-authentication processing memory, the authentication balance being updated to a balance amount after settlement when a transaction is settled by an authentication process having a requirement for a personal authentication to be matched, the non-authentication balance being... ..an amount less than or equal to the stored authentication balance when the transaction is settled by the authentication process; and

determining that in illegal process has been performed with the card when a comparison of the non-authentication balance and... ..than the authentication balance.

According to the present invention there is also provided a transaction settlement system comprising a transaction terminal and an electronic transaction card readable by said transaction terminal, wherein:

said card includes memory storing an authentication balance and a non-authentication balance, the authentication balance for transaction settlement by an authentication process requiring a personal authentication to be matched and the non-authentication balance for transaction settlement by a non-authentication process wherein the personal authentication is not required to be matched; and

said system further comprises a control unit controlling updates to the authentication balance and the non-authentication balance and detecting that an illegal process has been performed with said card when a comparison indicates the non-authentication balance is... ..embodiment of the present invention can provide an electronic cashing card which can execute a settlement operation by depositing an amount of money for personal authentication money processing by requesting matching of a personal authentication. Such a card can also maintain transaction security in non-personal authentication money processing wherein matching of the personal authentication is not requested.

An embodiment of the present invention also provides an electronic cashing card settlement system for executing settlement with a relevant electronic cashing card. The system comprises, within an electronic cashing card, a non-personal authentication money processing for memory for storing data in non-personal authentication money processing for executing settlement of money without a condition for matching of a personal authentication. A personal authentication money processing memory of the card stores data for personal authentication money

processing, whereby settlement of money with a condition for matching of a personal authentication can be executed. Thereby, when settlement is executed by the personal authentication money processing, the amount of an authentication balance to be stored in the balance area of the personal authentication money processing memory is updated by the amount of a balance after the settlement. Also, the amount of a non-authentication balance to be stored in the balance area of the non-personal authentication money processing memory is updated to an amount of money less than or equal to the authentication balance or a relevant amount of money. On the occasion of executing the settlement by the non-personal authentication money processing or the personal authentication money processing, the non-authentication balance and the authentication balance are compared. If the non-authentication balance... authentication balance as a result of the comparison, it is determined that a certain illegal process has been conducted with the electronic cashing card.

On the occasion of executing the settlement by the personal authentication money processing, whereby the settlement of money is executed under the condition of matching the personal authentication, the amount of money of the authentication balance of the personal authentication money processing memory is updated. Also, an amount of money of the non-authentication balance of the non-personal authentication money processing memory is updated to the amount of money equal to or smaller than the amount of the authentication balance. Moreover, on the occasion of executing the settlement by the non-personal authentication money processing, wherein settlement is executed without the condition of matching a personal authentication, it is determined that an illegal process has been executed with the electronic cashing card if the non-authentication balance is larger... result of comparison between these balance amounts.

Therefore, the money of the personal authentication money processing can be used in a disbursement process for the settlement by the non-personal authentication money processing. Moreover, transaction security can also be assured to improve convenience and safety of the card.

Furthermore, an embodiment of the present invention is characterized in that, when settlement is executed by the non-personal authentication money processing, a balance after the settlement is calculated on the basis of the non-authentication balance to be stored in the balance area of the non-personal authentication money processing memory. The non-authentication balance, in the non-personal authentication money processing memory, is updated to the balance after the settlement. When settlement is executed by the personal authentication money processing, the balance after settlement is calculated on the basis of the non-authentication balance stored in the balance area of the non-personal authentication money processing memory. The non-authentication balance of the non-personal authentication money processing memory and the authentication balance of the personal authentication money processing memory are updated to the balance after settlement.

Accordingly, when settlement is executed by the non-personal authentication money processing, it is no longer required to make access to the personal authentication money processing area. Therefore, security of the transaction can be assured. In the non-personal authentication money processing, a subtraction is conducted only by disbursing the non-authentication balance. Therefore, the non-authentication... a non-authentication balance exceeds the authentication balance, it is determined that a certain illegal process has been executed for the electronic cashing card.

In addition, when settlement is executed by the non-personal authentication processing of the invention, the authentication balance and the non-authentication balance are stored respectively in the personal authentication money processing area and non-personal authentication money processing area. Because these balances are distinct from each other, the balance after settlement is calculated on the basis of the non-authentication balance reflecting the result of settlement by the non-personal authentication money processing. When making settlement by the personal authentication money processing, synchronization can be taken by setting both the authentication balance and the non-authentication balance to the balance after the

settlement.

Moreover, when settlement is executed by the non-personal authentication money processing, the balance after the settlement is calculated on the basis of the non-authentication balance stored in the balance area of the non-personal authentication money processing memory. The non-authentication balance of the non-personal authentication money processing memory is updated to the balance after the settlement. When settlement is executed by the personal authentication money processing, the balance after the settlement is calculated on the basis of the authentication balance stored in the balance area of the personal authentication money processing memory and the non-personal authentication balance stored in the balance area of the non-personal authentication money processing memory. The authentication balance of the personal authentication money processing memory is updated to the balance after the settlement. Moreover, the non-authentication balance of the non-personal authentication money processing memory is updated to an amount of money smaller than the authentication balance under a predetermined condition.

Even in this structure, when the settlement is executed by the non-personal authentication money processing, settlement can be executed without making access to the personal authentication money processing area. This ensures transaction security, in that the non-authentication balance does not exceed the... ..the non-authentication balance exceeds the authentication balance, it is determined that a certain illegal process has been performed with the electronic cashing card.

On the other hand, in the synchronization processing between the authentication balance and non-authentication balance in the personal authentication money processing, these balances are never matched and the non-authentication balance is set to an amount... ..the authentication balance under a predetermined condition. Thereby, the amount of money available for a settlement by the non-personal authentication money processing is limited to assure greater transaction security.

When required, the settlement system of an embodiment of the present invention compares the non-authentication balance and the authentication balance for every predetermined number of transaction by the non-personal authentication money processing.

When required, the settlement system of an embodiment of the present invention sets a limit amount for settlements to be executed by the non-personal authentication money processing. The system determines that a certain illegal process has been executed with the electronic cashing card when the amount of a disbursement to be written as the disbursement history of the non-personal authentication money processing memory becomes larger than the settlement limit amount.

When required, when a deposit or settlement process is executed by the personal authentication money processing, the settlement system of an embodiment of the present invention deposits a predetermined amount or a predetermined rate to the non-personal authentication money processing in order to write such amount or rate to predetermined area of the non-personal authentication money processing memory.

When required, the settlement system of an embodiment of the present invention is provided, in the electronic cashing card... ..arithmetic means for executing various arithmetic calculations in regard to the non-personal authentication money processing and the personal authentication money processing. The arithmetic means also controls the read and write operations to the non-personal authentication and personal authentication money processing memories. Also provided as an input/output means for executing input and output of data... ..means and an external device.

An embodiment of the present invention further provides a transaction settlement system comprising a transaction terminal and an electronic transaction card readable by the terminal. The card includes memory storing an authentication balance and a non-authentication balance. The authentication balance is for transaction settlement by an authentication process requiring a personal authentication to be matched. The non-authentication balance is for

transaction settlement by a non-authentication process wherein the personal authentication is not required to be matched. The system further comprises a control unit controlling updates to the authentication balance and the non-authentication balance. The control unit detects that an illegal process has been performed with the card when a comparison indicates the non-authentication balance is... ..for example used for transactions involving larger amounts of money, time, or other tokens or units, it may also act as an authorisation for the card holder to access or use... ..with an embodiment of the present invention, no transaction based on the non-personal authentication processing aspect of the card settlement system accesses or opens a route for access to the personal authentication money processing aspect of the system. That is, the consumer or card user, having selected that a transaction should be based on the non-personal authentication processing aspect of the cashing card settlement system, is assured that this is the case and that no access can be made to the personal authentication processing aspect of the card settlement system. At the same time, however, the balance, or a part or proportion of the balance, in the personal authentication processing aspect of the card settlement system is made available for transactions based on the non-personal authentication money processing aspect of the cashing card settlement system. There is thus provided a satisfactory level of transaction security, together with convenient availability... ..amount which may be the subject of a transaction) - are made through the personal authentication processing aspect of the cashing card settlement system. This loaded balance may then be made available for settlement or "payment" transactions through the non-personal authentication processing aspect of the card settlement system. Alternatively, only a fixed, lesser amount may be made so available. Further alternatively, the... ..present invention can provide that an amount, for example of money, deposited for personal authentication processing can be used for a settlement process with the non-personal authentication processing, whilst ensuring transactional security and offering advantages of improved convenience and safety in the use...with the accompanying drawings of which:

FIG. 1 is a schematic block diagram of a settlement system of an electronic cashing card according to a preferred embodiment of the present invention;

FIG. 2 is a data layout format in the settlement system illustrated in FIG. 1;

FIG. 3 illustrates a format for each of disbursement data, deposit data, balance data, and protection limitation data of non-personal authentication money processing in the data layout illustrated in FIG. 2;

FIG. 4 illustrates a schematic operation flowchart for a total settlement process of the settlement system illustrated in FIG. 1;

FIG. 5 illustrates a detail operation flowchart of step SP... ..the total schematic operation flowchart illustrated in FIG. 4;

FIG. 7 illustrates operations in the synchronization process illustrated in FIG. 6;

FIG. 8 illustrates an operation flowchart of a deposit process to the non-personal authentication money processing area in a settlement system according to a second embodiment of the present invention;

FIG. 9A is a schematic... ..IC card type electronic cashing card illustrated in FIG. 9A; and

FIG. 11 illustrates a settlement operation flowchart in the non-personal authentication and personal authentication money processing of the electronic cashing card of the related system illustrated in FIG. 9A.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

The settlement system of the present invention will now be explained, with reference in detail to presently... ..illustrated in the accompanying drawings, wherein like reference numerals refer to like elements throughout.

The settlement system of a first embodiment of the present invention will be explained with reference to FIGS. 1-4. FIG. 1 is a schematic block diagram of the settlement system of this embodiment. FIG. 2 is a diagram illustrating a data layout format in the settlement system illustrated in FIG. 1. FIG. 3 is a diagram illustrating a format for each...

...data, deposit data, balance data, and protection limit data in the non-personal authentication money processing with the data layout illustrated in FIG. 2.

FIGS. 4-6 each illustrate an operation flowchart of the settlement process of the electronic cashing card settlement system illustrated in FIG. 1.

Like the electronic cashing card in the system of the related art illustrated in FIG. 9, the electronic cashing card settlement system of this embodiment, as illustrated in each figure, comprises a write/read control means... 4, and an input/output means 7. Also included in this embodiment is a money processing memory 1 for each data of the non-personal authentication money processing (for executing settlement of money without request for a password number as a personal authentication), the personal authentication money processing (for executing settlement of money with a request for a password number as a personal authentication) and the management items in common for these money processes, as explained above.

Moreover, a synchronous processing means 5 executes synchronization processing under a predetermined condition for each balance of the non-personal authentication money processing and the personal authentication money processing, the balances being stored in the money processing memory 1. This synchronization processing is executed under the control of the control arithmetic means 4. A comparing arithmetic means 6 compares balances of the non-personal authentication money

processing and the personal authentication money processing stored in the money processing memory 1, under the control of the control arithmetic means 4. Therefore, as a result... arithmetic comparison by the comparing arithmetic means 6, it is determined that a certain illegal process has been conducted with the IC card 10 when the balance of the non-personal authentication money processing is determined to be larger.

As shown in FIG. 2, the money processing memory 1 is structured to comprise a common area 11 for storing the data common to the non-personal authentication money processing and personal authentication money processing in regard to the IC card 10. A non-personal authentication money processing area 12 stores data of non-personal authentication money processing, and a personal authentication money processing area 13 stores data of personal authentication money processing.

The common area 11 includes a regional code of a management organization, a company code... personal information includes a password number and time limiting information. The non-personal authentication money processing area 12 includes a disbursement content history area 12a that stores as history data a disbursement content, obtained by execution of the settlement process, of the amount of sales by the non-personal authentication money processing. The history data is stored in the data format illustrated in FIG. 3A. A balance area 12b stores a non-authentication balance used by the non-personal authentication money process after execution of the settlement process for the amount of sales. The non-authentication balance is stored in the data format... information area 12c stores various pieces of limitation information for executing non-personal authentication money processing. The limitation information is stored in the data format illustrated in FIG. 3D.

This limitation information area 12c stores various data, such as the number of times of continuous processes as accumulated data, i.e., a count of the number of continuous executions of the settlement process by the non-personal authentication money processing. Also stored is a continuous total amount obtained by totaling the amounts of sale settled by such continuous processes. A continuous process upper limit value limits the allowed continuous processes, and a continuous amount upper limit value limits the continuous total amount for the continuous processes. The number of times of continuous processes and the continuous total amount are sequentially added for each execution of a continuous settlement process by the non-personal authentication money processing. When a settlement process or a deposit process is executed by the personal authentication money processing, such values are reset to "0" respectively.

The personal authentication money processing area 13 includes a disbursement content history area 13a for storing as history data a disbursement content, obtained by executing the settlement process, for an amount of sales by the personal authentication money processing. The history (i.e., disbursement) data is stored in the data format illustrated in FIG. 3A. Area 13b stores an authentication balance amount to be used by the personal authentication money processing after the settlement process for amount of sales is executed. The authentication balance amount is stored in the data format illustrated in FIG. 3B. A deposit content history area 13c stores the deposit content for the personal authentication money processing in the data format illustrated in FIG. 3B.

Next, operation of the settlement process executed by the electronic cashing card settlement system of the present embodiment, based on the structure explained above, will be explained with reference to FIG. 4.

In FIG. 4, it is determined first whether the process by a consumer (as the owner of the IC card 10) should execute a disbursement operation or a deposit operation using the IC card 10 (step SP 1). When the disbursement process is determined to be executed in step SP 1, an amount of sales is input (step SP 2). A password number is determined to be not input in step SP 3, a series of processes for the settlement process are executed by the non-personal authentication money processing (step SP 4). When the password number is determined to be input in step SP 3, a series of processes for a settlement process are executed by the personal authentication money processing (step SP 6).

Moreover, when the deposit operation is determined to be executed in step SP 1, a series of operations for a deposit process are executed by the personal authentication money processing (step SP 7).

The settlement process by the non-personal authentication money processing in step SP 4 of FIG. 4 is illustrated by the flowchart of FIG. 5. A non-authentication balance of the non-personal authentication money processing area 12 is read from the balance area 12b (step SP 41). Moreover, the non-authentication balance of the non-personal authentication money processing area 12 is read from the balance area 12b (step SP 42). In addition, the authentication balance of the personal authentication money processing area 13 is read from the balance area 13b (step SP 43). This read non-personal balance is compared with the authentication balance in the comparing arithmetic unit 6 (step SP 44).

When the non-authentication balance is determined to be the larger value in step SP 44, a transaction cease process is executed under the assumption that an illegal process has been performed with the IC card 10 (step SP 45). Meanwhile, in step SP 46, or the values are equal, "1" is added to the number of times of continuous processes. Thereby an accumulated value of the number of times of continuous execution of the settlement process by the non-personal authentication money processing is indicated (step SP 46). Moreover, in step SP 46, the amount of sales can also be added to the continuous total amount, which is a total amount of the settlement processes of the number of times of continuous processes. In step SP 46.

It is then determined whether the added number of times of continuous processes is a value equal to or smaller than the upper limit value of continuous processes stored in the limitation information area 12c (step SP 47). Whether the continuous total amount is determined to be equal to or smaller than the upper limit of continuous processes in step SP 47, or equal to or smaller than the upper limit value of the continuous total amount, the settlement process by the non-personal authentication money processing is executed (step SP 48). The non-authentication balance calculated by this settlement process updates the non-authentication balance of the balance area 12b (step SP 49) and is... a new non-authentication balance (step SP 50). Moreover, a disbursement content, based on the settlement process executed in step SP 48, is written to the disbursement content history area 12a (step SP 50).

When the number of times of continuous processes is determined to be equal to or larger than the upper limit of continuous processes in step SP 47, or when the continuous total amount is determined to be equal to or larger than the upper limit value of the continuous amount, a message indicating that settlement by the non-personal authentication money processing is disabled is notified or displayed from a register or a terminal equipment or the like of the POS system (step SP 55). For example, as a notifying and displaying method, settlement by personal authentication money processing or settlement by cash can be selected.

FIG. 6A illustrates the settlement process by the personal authentication processing in step SP 6. First, It is determined whether the password number matches (step SP 61). The settlement process by the personal authentication money processing is executed under the condition that the passwords are matched (step SP 62). On the basis of the authentication balance after this settlement process, the synchronous processing means 5 executes the synchronization process (step SP 63). This synchronization process updates the disbursement content history area 12a of the non-personal authentication money processing area 12, under the condition that the settlement process is executed by the personal authentication money processing. The disbursement content history area 13a of the personal authentication money processing area 13 is also updated. Further, the authentication balance and the non-authentication balance are updated. After this synchronization process is executed, each value is reset to "0" by resetting the accumulated values for the number of times of continuous processes and continuous total amount of money (step SP 64).

In the case of a deposit process by the personal authentication money processing in step SP 7, the password number is input (step SP 71) and then the deposit process is executed by cash or from the account with the condition of matching of the password (step SP 73). Thereafter, the synchronous processing means 5 performs the synchronization process based on the authentication balance after this deposit process (step SP 74). This synchronization process is performed by updating the disbursement content history area 12a of the non-personal authentication money processing area 12. The disbursement content history area 13a of the personal authentication money processing area 13 is also updated as required, and the authentication balance and non-authentication balance are updated. In this case, the updates are executed under the condition that the deposit process has been executed by the personal authentication money processing. After this synchronization process has been executed, the accumulated value of the number of times of continuous processes and the continuous total amount are reset to initialize each value to "0" (step SP 75).

The practical synchronization process will be explained with reference to FIGS. 7A and 7B. First, a practical example will be described with reference to FIG. 7A. When a deposit process for (Yen)10,000 (

process 1) is executed by the personal authentication money processing, (Yen)10,000 is written to the deposit content history area 13c and the balance... ..10,000 is written to the balance area 12b of the non-personal authentication money processing area 12. When a settlement process (process 2) for (Yen)2,000 is executed by the non-personal authentication money processing in such a deposit condition, only the disbursement content history area 12a and the balance area 12b are updated to (Yen)8,000.

Next, when a deposit process (process 3) for (Yen)20,000 is executed by the personal authentication money processing, (Yen)20,000 is written to the deposit content history area 13c, as in process 1. The authentication balance of the balance area 13b and the authentication balance of the... ..are also updated to (Yen)28,000.

The balance of the non-personal authentication money processing area 12 reflects the disbursement process of the non-personal authentication money processing by process 2, while the balance of the personal authentication money processing area 13 does not reflect such a disbursement process. Therefore, the balance is calculated on the basis of the balance of the non-personal authentication money processing area 12 and the deposit process amount, and both balances of the balance areas 12b, 13b are updated. Thereby, the contents of the non-personal authentication money processing area 12 and the personal authentication money processing area 13 can be synchronized.

Moreover, when a settlement process (process 4) for (Yen)4,000 is performed by the non-personal authentication money processing, (Yen)4,000 is written to the disbursement content history area 12a as in the case of the process 2, and the balance area 12b is updated to (Yen)24,000.

In addition, when a settlement process (process 5) for (Yen)15,000 is performed by the personal authentication money processing, (Yen)15,000 is written to the disbursement content history area 13a, as in the case of the deposit

process of the processes 1 and 3. Balance area 13b and balance area 12b are updated to (Yen)9,000. This update is performed according to the balance in the non-personal authentication money processing area 12 and the balance calculated on the basis of the settlement process amount. Thereby, the contents of the non-personal authentication money processing area 12 and personal authentication money process area 13 may be synchronized.

For settlement by the non-personal authentication money processing, only the disbursement content history area 12a and the balance 12b of the non-personal authentication money processing area 12 are updated. The synchronization process is executed in this case for each deposit or transaction settlement by the personal authentication money processing. Accordingly, in a regular application of the IC card 10, the non-authentication balance of the non-personal authentication money processing area 12 does not exceed the authentication balance of the personal authentication money processing area 13. Therefore, an illegal process with the IC card 10 can be determined by comparing the amounts of the non... ..balance and the authentication balance of the balance areas 12b, 13b through execution of the synchronization process.

In FIG. 7B, the non-authentication balance of the non-personal authentication money processing area 12 is limited to (Yen)5,000 or less. In this case, when a deposit process of (Yen)10,000 is performed by the personal authentication money processing, the deposit amount to the balance area 12b of the non-personal authentication money processing area 12 (written as an additional amount) is limited to (Yen)5,000. Limitation to... ..predetermined rate and the desired amount of money of (Yen)10,000 for the deposit process, in addition to the amount of (Yen)5,000. The other operations will be executed as in the FIG. 7A case.

FIG. 8 illustrates an operation flowchart of a deposit process for non-personal authentication money processing in a system according to a second embodiment of the present invention.

In FIG. 8, the electronic cashing card settlement system of this embodiment is provided with a structure to detect execution of an illegal process with respect to the non-personal authentication money processing area 12. This detection is performed by comparison of the non-authentication balance and the... ..system of FIGS 1-7. The structure of this second embodiment also executes a disbursement process to the non-personal authentication money processing area 12 from the personal authentication money processing area 13.

The electronic cashing card settlement system of this embodiment, based on the structure explained above, will described with reference to... ..the amount of sales in step SP 51, an amount from the personal authentication money processing is disbursed to the non-personal authentication money processing side, but only in a preset amount (step SP 52). This disbursement process can have its limit amount for disbursement set to an amount of money predetermined depending on the owner of IC card 10 or the application area (class of shopping stores) of IC card 10; or to a predetermined amount or an amount of a predetermined rate for each deposit and disbursement by the personal authentication money processing; or to an amount limiting the non-authentication balance to a predetermined amount or less.

When the predetermined amount of money is deposited by the disbursement process to the non-personal authentication money processing area, the non-authentication balance is updated by the new non-authentication balance. The new... ..authentication balance and the authentication balance, determinations for whether the upper limit value of continuous processes or the upper limit value of continuous amount, and arithmetic operations of various kinds may... ..operations may be executed by the controlled arithmetic means 4, write controlling means 3, synchronous processing means 5 and comparison arithmetic means 6 of the IC card 10. Alternatively, such determinations and operations may be executed by a terminal unit having a R/W unit, an ATM, or a POS device installed on the shopping store side. Moreover, such determinations and operations may also be executed by stored processes between each of the arithmetic means explained above.

In the present invention, a settlement process is executed by a personal authentication money processing to execute settlement of money under the condition of matching a personal authentication. A personal authentication balance of the personal authentication money processing memory is updated and a non-authentication balance of the non-

personal authentication money processing memory is updated to an amount equal to or smaller than the authentication balance.

It... ..balance. This may occur as a result of comparison on the occasion of executing a settlement process by a non-personal authentication money processing which executes settlement of money without a condition for matching the personal authentication. In this situation it is determined that an illegal process has been performed with the electronic cashing card.

These features allow an amount of money from the personal authentication money processing to be deposited for use by a settlement process with the non-personal authentication money processing. Also, transaction security can be assured, providing the advantages of greatly improved convenience and safety...

Specification: ...authentication money processing in regard to the IC card 10. A non-personal authentication money processing area 12 stores data of non-personal authentication money processing, and a personal authentication money processing area 13 stores data of personal authentication money processing.

The common area 11 includes a regional code of a management organization, a company code... ..personal information includes a password number and time limiting information. The non-personal authentication money processing area 12 includes a disbursement content history area 12a that stores as history data a disbursement content, obtained by execution of the settlement process, of the amount of sales by the non-personal authentication money processing. The history data is stored in the data format illustrated in FIG. 3A. A balance area 12b stores a non-authentication balance used by the non-personal authentication money process after execution of the settlement process for the amount of sales. The balance is stored in the data format illustrated in... ..information area 12c stores various pieces of limitation information for executing non-personal authentication money processing. The limitation information is stored in the data format illustrated in FIG. 3D.

This limitation information area 12c stores various data, such as the number of times of continuous processes as accumulated data, i.e., a count of the number of continuous executions of the settlement process by the non-personal authentication money processing. Also stored is a continuous total amount obtained by totaling the amounts of sale settled by such continuous processes. A continuous process upper limit value limits the allowed continuous processes, and a continuous amount upper limit value limits the continuous total amount for the continuous processes. The number of times of continuous processes and the continuous total amount are sequentially added for each execution of a continuous settlement process by the non-personal authentication money processing. When a settlement process or a deposit process is executed by the personal authentication money processing, such values are reset to "0" respectively.

The personal authentication money processing area 13 includes a disbursement content history area 13a for storing as history data a disbursement content, obtained by executing the settlement process, for an amount of sales by the personal authentication money processing. The history (i.e., disbursement) data is stored in the data format illustrated in FIG... ..area 13b stores an authentication balance amount to be used by the personal authentication money

processing after the settlement process for amount of sales is executed. The authentication balance amount is stored in the data... ..A deposit content history area 13c stores the deposit content for the personal authentication money processing in the data format illustrated in FIG. 3B.

Next, operation of the settlement process executed by the electronic cashing card settlement system of the present embodiment, based on the structure explained above, will be explained with reference to FIG. 4.

In FIG. 4, it is determined first whether the process by a consumer (as the owner of the IC card 10) should execute a disbursement operation or a deposit operation using the IC card 10 (step SP 1). When the disbursement process is

determined to be executed in step SP 1, an amount of sales is input... ..password number is determined to be not input in step SP 3, a series of processes for the settlement process are executed by the non-personal authentication money processing (step SP 4). When the password number is determined to be input in step SP 3, a series of processes for a settlement process are executed by the personal authentication money processing (step SP 6).

Moreover, when the deposit operation is determined to be executed in step SP 1, a series for operations for a deposit process are executed by the personal authentication money processing (step SP 7).

The settlement process by the non-personal authentication money processing in step SP 4 of FIG. 4 is illustrated by the flowchart of FIG. 5... ..read (step SP 41). Moreover, the non-authentication balance of the non-personal authentication money processing area 12 is read from the balance area 12b (step SP 42). In addition, the authentication balance of the personal authentication money processing area 13 is read from the balance area 13b (step SP 43). This read non-personal balance is compared with the authentication balance in the comparing arithmetic unit 6 (step SP 44).

When the non-authentication balance is determined to be the larger value in step SP 44, a transaction cease process is executed under the assumption that an illegal process has been performed with the IC card 10 (step SP 45).

Meanwhile, in step SP... ..or the values are equal, "1" is added to the number of times of continuous processes.

Thereby an accumulated value of the number of times of continuous execution of the settlement process by the non-personal authentication money processing is indicated (step SP 46). Moreover, in step SP 46, the amount of sales can also be added to the continuous total amount, which is a total amount of the settlement processes of the number of times of continuous processes In step SP 46.

It is then determined whether the added number of times of continuous processes is a value equal to or smaller than the upper limit value of continuous processes stored in the limitation information area 12c (step SP 47). Whether the continuous total amount... ..amount is determined to be equal to or smaller than the upper limit of continuous processes in step SP 47, or equal to or smaller than the upper limit value of the continuous total amount, the settlement process by the non-personal authentication money processing is executed (step SP 48). The non-authentication balance calculated by this settlement process updates the non-authentication balance of the balance area 12b (step SP 49) and is... ..a new non-authentication balance (step SP 50). Moreover, a disbursement content, based on the settlement process executed in step SP 48, is written to the disbursement content history area 12a (step SP 50).

When the number of times of continuous processes is determined to be equal to or larger than the upper limit of continuous processes in step SP 47, or when the continuous total amount is determined to be equal... ..or larger than the upper limit value of the continuous amount, a message indicating that settlement by the non-personal authentication money processing is disabled is notified or displayed from a register or a terminal equipment or the like of the POS system (step SP 55). For example, as a notifying and displaying method, settlement by personal authentication money processing or settlement by cash can be selected.

FIG. 6A illustrates the settlement process by the personal processing in step SP 6. First, it is determined whether the password number matches (step SP 61). The settlement process by the personal authentication money processing is executed under the condition that the passwords are matched (step SP 62). On the basis of the authentication balance after this settlement process, the synchronous processing means 5 executes the synchronization process (step SP 63). This synchronization process updates the disbursement content history area 12a of the non-personal authentication money processing area 12, under the condition that the settlement process is executed by the personal money processing. The disbursement content history area 13a of the personal authentication money processing area 13 is also updated. Further, the authentication balance and the non-authentication balance are updated. After this synchronization process is executed, each value is reset to "0" by resetting the accumulated values for the number of times of continuous processes and continuous total amount of money (step SP 64).

In the case of a deposit process by the personal authentication money processing in step SP 7, the password number is input (step SP 71) and then the deposit process is executed by cash or from the account with the condition of matching of the password (step SP 73). Thereafter, the synchronous processing means 5 performs the synchronization process based on the balance after this deposit process (step SP 74). This synchronization process is performed by updating the disbursement content history area 12a of the non-personal money processing area 12. The disbursement content history area 13a of the personal authentication money processing area 13 is also updated as required, and the authentication balance and non-authentication balance are updated. In this case, the updates are executed under the condition that the deposit process has been executed by the personal authentication money processing. After this synchronization process has been executed, the accumulated value of the number of times of continuous processes and the continuous total amount are reset to initialize each value to "0" (step SP 75).

The practical synchronization process will be explained with reference to FIGS. 7A and 7B. First, a practical example will be described with reference to FIG. 7A. When a deposit process for (Yen)10,000 (process 1) is executed by the personal authentication money processing, (Yen)10,000 is written to the deposit content history area 13c and the balance... ..10,000 is written to the balance area 12b of the non-personal authentication money processing area 12. When a settlement process (process 2) for (Yen)2,000 is executed by the non-personal authentication money processing in such a deposit condition, only the disbursement content history area 12a and the balance area 12b are updated to (Yen)8,000.

Next, when a deposit process (process 3) for (Yen)20,000 is executed by the personal authentication money processing, (Yen)20,000 is written to the deposit content history area 13c, as in process 1. The balance of the balance area 13b and the authentication balance of the balance... ..are also updated to (Yen)28,000.

The balance of the non-personal authentication money processing area 12 reflects the disbursement process of the non-personal authentication money processing by process 2, while the balance of the personal authentication money processing area 13 does not reflect such a disbursement process. Therefore, the balance is calculated on the basis of the balance of the non-personal authentication money processing area 12 and the deposit process amount, and both balances of the balance areas 12b, 13b are updated. Thereby, the contents of the non-personal authentication money processing area 12 and the personal authentication money processing area 13 can be synchronized.

Moreover, when a settlement process (process 4) for (Yen)4,000 is performed by the non-personal authentication money processing, (Yen)4,000 is written to the disbursement content history area 12a as in the case of the process 2, and the balance area 12b is updated to (Yen)24,000.

In addition, when a settlement process (process 5) for (Yen)15,000 is performed by the personal authentication money processing, (Yen)15,000 is written to the disbursement content history area 13a, as in the case of the deposit process of the processes 1 and 3. Balance area 13b and balance area 12b are updated to (Yen)9,000. This update is performed according to the balance in the non-personal authentication money processing area 12 and the balance calculated on the basis of the settlement process amount. Thereby, the contents of the non-personal authentication money processing area 12 and personal authentication money processing area 13 may be synchronized.

For settlement by the non-personal authentication money processing, only the disbursement content history area 12a and the balance 12b of the non-personal authentication money processing area 12 are updated. The synchronization process is executed in this case for each deposit or transaction settlement by the personal authentication money processing. Accordingly, in a regular application of the IC card 10, the non-authentication balance of the non-personal authentication money processing area 12 does not exceed the authentication balance of the personal authentication money processing area 13. Therefore, an illegal process with the IC card 10 can be determined by comparing the amounts of the non... ..balance and the authentication balance of the balance areas 12b, 13b through execution of the synchronization process.

In FIG. 7B, the non-authentication balance of the non-personal authentication money

processing area 12 is limited to (Yen)5,000 or less. In this case, when a deposit process of (Yen)10,000 is performed by the personal authentication money processing, the deposit amount to the balance area 12b of the non-personal authentication money processing area 12 (written as an additional amount) is limited to (Yen)5,000. Limitation to... ..predetermined rate and the desired amount of money of (Yen)10,000 for the deposit process, in addition to the amount of (Yen)5,000. The other operations will be executed as in the FIG. 7A case.

FIG. 8 illustrates an operation flowchart of a deposit process for non-personal authentication money processing in a system according to a second embodiment of the present invention.

In FIG. 8, the electronic cashing card settlement system of this embodiment is provided with a structure to detect execution of an illegal process with respect to the non-personal authentication money processing area 12. This detection is performed by comparison of the non-authentication balance and the... ..system of FIGS 1-7. The structure of this second embodiment also executes a disbursement process to the non-personal authentication money processing area 12 from the personal authentication money processing area 13.

The electronic cashing card settlement system of this embodiment, based on the structure explained above, will be described with reference to... ..the amount of sales in step SP 51, an amount from the personal authentication money processing is disbursed to the non-personal authentication money processing side, but only in a preset amount (step SP 52). This disbursement process can have its limit amount for disbursement set to an amount of money predetermined depending on the owner of IC card 10 or the application area (class of shopping stores) of IC card 10; or to a predetermined amount or an amount of a predetermined rate for each deposit and disbursement by the personal authentication money processing; or to an amount limiting the non-authentication balance to a predetermined amount or less.

When the predetermined amount of money is deposited by the disbursement process to the non-personal authentication money processing area, the non-authentication balance is updated by the new non-authentication balance. The new... ..authentication balance and the authentication balance, determinations for whether the upper limit value of continuous processes or the upper limit value of continuous amount, and arithmetic operations of various kinds may... ..operations may be executed by the controlled arithmetic means 4, write controlling means 3, synchronous processing means 5 and comparison arithmetic means 6 of the IC card 10. Alternatively, such determinations and operations may be executed by a terminal unit having a R/W unit, an ATM, or a POS device installed on the shopping store side. Moreover, such determinations and operations may also be executed by shared processes between each of the arithmetic means explained above.

In the present invention, a settlement process is executed by a personal authentication money processing to execute settlement of money under the condition of matching a personal authentication. A personal authentication balance of the personal authentication money processing memory is updated and a non-authentication balance of the non-personal authentication money processing memory is updated to an amount equal to or smaller than the authentication balance.

It... ..balance. This may occur as a result of comparison on the occasion of executing a settlement process by a non-personal authentication money processing which executes settlement of money without a condition for matching the personal authentication. In this situation it is determined that an illegal process has been performed with the electronic cashing card.

These features allow an amount of money from the personal authentication money processing to be deposited for use by a settlement process with the non-personal authentication money processing. Also,

Claims:

1. A system for transaction settlement with an electronic cashing card having a non-authentication processing memory and an authentication processing memory, said system comprising:

means for updating an authentication balance stored in a balance area of the authentication processing memory and a non-authentication balance stored in a balance area of the non-authentication processing memory, said means updating the authentication balance to a balance amount after settlement when a transaction is settled by an authentication process having a requirement for a personal authentication to be matched, said means updating the non... ..an amount less than or equal to the stored authentication balance when the transaction is settled by the authentication process; and

means for comparing the non-authentication balance and the authentication balance and determining that an illegal process has been performed with the card when the non-authentication balance is larger than the authentication balance.

2. A system as claimed in claim 1, wherein:

when a transaction is settled by a non-authentication process wherein the personal authentication is not required to be matched, the balance amount after settlement is calculated based on the stored authentication balance and the non-authentication balance is updated to the balance amount after settlement; and

when a transaction is settled by the authentication process, the balance amount after settlement is calculated based on the stored non-authentication balance and both the authentication balance and the non-authentication balance are updated to the balance amount after settlement.

3. A system as claimed in claim 1, wherein:

when a transaction is settled by a non-authentication process wherein the personal authentication is not required to be matched, the balance amount after settlement is calculated based on the stored non-authentication balance and the non-authentication balance is updated to the balance amount after settlement; and

when a transaction is settled by the authentication process, the balance amount after settlement is calculated based on the stored authentication balance and the stored non-authentication balance, the authentication balance is updated to the balance amount after settlement, and the non-authentication balance is updated according to a preset condition amount.

4. A... ..balance and the non-authentication balance are compared in each of successive transactions to be settled by a non-authentication process wherein the personal authentication is not required to be matched, when a count of the... ..number.

5. A system as claimed in claim 1, 2, 3 or 4, wherein:

a settlement amount limit is set for settlement of transactions by a non-authentication process wherein the personal authentication is not required to be matched; and

said comparing and determining means determines that an illegal process has been performed with the card when a disbursement amount, to be written in the non-authentication processing memory as a disbursement history, exceeds the settlement amount limit.

6. A system as claimed in any preceding claim, wherein when the authentication process is invoked to perform a deposit or to settle a transaction, a money amount is deposited for the authentication process and is written to a predetermined area of the non-authentication processing memory, the money amount comprising at least one of a predetermined cash amount and a... ..claim, further comprising within the card:

arithmetic means for executing arithmetic calculations for the authentication process and a non-authentication process wherein the personal authentication is not required to be matched, said arithmetic means further controlling

data reading and writing operations from and to the non-authentication processing memory and the authentication processing memory; and

input/output means for executing data input/output operations between the arithmetic means and an external unit.

8. A system as claimed in any preceding claim, wherein the card comprises an integrated... any preceding claim, wherein the card is a prepaid card.

10. A method of transaction settlement with an electronic cashing card having a non-authentication processing memory and an authentication processing memory, the method comprising:

updating an authentication balance stored in the authentication processing memory and a non-authentication balance stored in the non-authentication processing memory, the authentication balance being updated to a balance amount after settlement when a transaction is settled by an authentication process having a requirement for a personal authentication to be matched, the non-authentication balance being... an amount less than or equal to the stored authentication balance when the transaction is settled by the authentication process; and

determining that an illegal process has been performed with the card when a comparison of the non-authentication balance and... authentication balance.

11. A method as claimed in claim 10, wherein:

when a transaction is settled by a non-authentication process wherein the personal authentication is not required to be matched, the balance amount after settlement is calculated based on the stored authentication balance and the non-authentication balance is updated to the balance amount after settlement; and

when a transaction is settled by the authentication process, the balance amount after settlement is calculated based on the stored non-authentication balance and both the authentication balance and the non-authentication balance are updated to the balance amount after settlement.

12. A method as claimed in claim 10, wherein:

when the transaction is settled by a non-authentication process wherein the personal authentication is not required to be matched, the balance amount after settlement is calculated based on the stored non-authentication balance and the non-authentication balance is updated to the balance amount after

settlement; and

when a transaction is settled by the authentication process, the balance amount after settlement is calculated based on the stored authentication balance and the stored non-authentication balance, the authentication balance is updated to the balance amount after settlement, and the non-authentication balance is updated to a preset condition amount.

13. A method... balance and the non-authentication balance are compared in each of successive transactions to be settled by a non-authentication process wherein the personal authentication is not required to be matched, when a count of the... as claimed in claim 10, 11, 12 or 13, further comprising determining that an illegal process has been performed with the card when a settlement amount limit is less than a disbursement amount to be written in the non-authentication processing memory as a disbursement history, the settlement amount limit being set for settlement of transactions by a non-authentication process wherein the personal authentication is not required to be matched.

15. A method as claimed in any of claims 10 to 14, wherein the authentication process is invoked to perform a deposit or to settle a transaction, a money amount is deposited for the authentication process and is written to a predetermined area of the non-authentication processing memory, the money amount comprising at least one of a predetermined cash amount and a... claims 10 to 15, further comprising:

executing within the card arithmetic calculations for the authentication process and anon-authentication process wherein the personal authentication is not required to be matched;

controlling within the card data reading and writing operations from and to the non-authentication processing memory and the authentication processing memory; and

executing within the card data input/output operations between the card and an external unit.

17. A computer program for settlement of transactions with an electronic cashing card having a non-authentication processing memory and an authentication processing memory, said program comprising procedures for:

updating an authentication balance stored in the authentication processing memory and a non-authentication balance stored in the non-authentication processing memory, the authentication balance being updated to a balance amount after settlement when a transaction is settled by an authentication process having a requirement for a personal authentication to be matched, the non-authentication balance being... ..an amount less than or equal to the stored authentication balance when the transaction is settled by the authentication process; and

determining that an illegal process has been performed with the card when a comparison of the non-authentication balance and... ..balance indicates that the non-authentication balance is larger than the authentication balance.

18. A computer readable medium encoded with the program of claim 17.

19. A computer program or a computer readable medium, as the case may be, as claimed in claim 17 or 18, wherein:

when a transaction is settled by a non-authentication process wherein the personal authentication is not required to be matched, the balance amount after settlement is calculated based on the stored authentication balance and the non-authentication balance is updated to the balance amount after settlement; and

when a transaction is settled by the authentication process, the balance amount after settlement is calculated based on the stored non-authentication balance and both the authentication balance and the non-authentication balance are updated to the balance amount after settlement.

20. A computer program or a computer readable medium, as the case may be, as claimed in claim 17 or 18, wherein:

when a transaction is settled by a non-authentication process wherein the personal authentication is not required to be matched, the balance amount after settlement is calculated based on the stored non-authentication balance and the non-authentication balance is updated to the balance amount after settlement; and

when a transaction is settled by the authentication process, the balance amount after settlement is calculated based on the stored authentication balance and the stored non-authentication balance, the authentication balance is updated to the balance amount after settlement, and the authentication balance is updated to a preset condition amount.

21. A computer program or a computer readable medium, as the case may be, as claimed in claim 17, 18, 19 or... ..balance and the non-authentication balance are compared in each of successive transactions to be settled by a non-authentication process wherein the personal authentication is not required to be matched, when a count of the successive transactions is less than or equal to a predetermined number.

22. A computer program or a computer readable medium, as the case may be, as claimed in any of claims 17 to 21, wherein said program further comprises a procedure for determining that an illegal process has been performed with the card when a settlement amount limit is less than a disbursement amount to be written in the non-authentication processing memory as a disbursement history, the settlement amount limit being set for settlement of transactions by a non-authentication process wherein the personal authentication is not required to be matched.

23. A computer program or a computer readable medium, as the case may be, as claimed in any of claims 17 to 22, wherein said program further comprises procedures for depositing a money amount for the authentication process and writing the money amount to a predetermined area of the non-authentication processing memory when the authentication process is invoked to perform a deposit or to settle a transaction, the money amount comprising at least one of a predetermined cash amount and a predetermined rate amount.

24. A computer program or a computer readable medium, as the case may be, as claimed in any of claims 17 to...
...said program further comprises procedures for:

executing within the card arithmetic calculations for the authentication process and a non-authentication process wherein the personal authentication is not required to be matched;

controlling within the card data reading and writing operations from and to the non-authentication processing memory and the authentication processing memory; and

executing within the card data input/output operations between the card and an external unit.

25. A transaction settlement system comprising a transaction terminal and an electronic transaction card readable by said transaction terminal, wherein:

said card includes memory storing an authentication balance and a non-authentication balance, the authentication balance for transaction settlement by an authentication process requiring a personal authentication to be matched and the non-authentication balance for transaction settlement by a non-authentication process wherein the personal authentication is not required to be matched; and

said system further comprises a control unit controlling updates to the authentication balance and the non-authentication balance and detecting that an illegal process has been performed with said card when a comparison indicates the non-authentication balance is...
...26. A system as claimed in claim 25, wherein said system further comprises a comparison unit comparing the stored authentication balance and the stored non-authentication balance when a transaction is to be settled with said card.

27. A system as claimed in claim 25 or 26, wherein an authentication processing memory and a non-authentication processing memory are included in said memory, the authentication processing memory storing the authentication balance and the non-authentication memory storing the non-authentication balance.

28. A system as claimed in claim 25, 26 or 27, wherein a money processing memory storing both the authentication balance and the non-authentication balance is included in the...
...claimed in any of claims 25 to 29, wherein said system further comprises an arithmetic unit for performing arithmetic calculations for the authentication process and the non-authentication process.

31. A system as claimed in any of claims 25 to 30, wherein said system further comprises a synchronization unit executing a synchronization process after a deposit has been performed by the authentication process, the synchronization process including updating a non-authentication disbursement history in the memory based on the authentication balance...
...controller updating the authentication balance and the non-authentication balance under control of the control unit.

33. A system as claimed in any of claims 25 to 32, wherein said control unit is included in said card.

34. A system as claimed in any of claims 25 to 33, wherein said comparison unit is included in said card.

35. A system as claimed in any of claims 25 to 34, wherein said arithmetic unit is included in said card.

36. A system as claimed in any of claims 25 to 35, wherein said synchronization unit is included in said card.

37. The card of the system as claimed in any...

Claims: ...der Karte und einer externen Einheit.

25. System gemas Anspruch 1, ferner mit einem Transaktions-Terminal, mit dem die Karte lesbar ist.

26. System gemas Anspruch 25, bei dem ein Geldverarbeitungsspeicher...

Claims: ...la base du montant du bilan apres reglement lorsqu'une transaction est reglee par un processus d'authentification necessitant qu'une correspondance d'authentification personnelle soit etablie, et mettant a jour...
...inferieur ou egal au bilan des authentifications memorise lorsque la transaction est reglee par le processus d'authentification ; et ledit moyen mettant a jour uniquement le bilan des non-authentifications sur...
...base d'un montant du bilan apres reglement lorsque la transaction est reglee par un processus

de non-authentification,

un moyen destine a comparer le bilan des non-authentifications et le bilan des authentifications et a determiner qu'un processus illegal a ete execute avec la carte lorsque le bilan des non-authentifications est superieur...
...2. Systeme selon la revendication 1, dans lequel :

lorsqu'une transaction est reglee par un processus de non-authentification dans lequel une correspondance a l'authentification personnelle n'est pas requise...
...base du montant du bilan apres reglement ; et

lorsqu'une transaction est reglee par le processus d'authentification, le montant du bilan apres reglement est calcule sur la base du bilan...
...3. Systeme selon la revendication 1, dans lequel :

lorsqu'une transaction est reglee par un processus de non-authentification dans lequel la correspondance a l'authentification personnelle n'est pas requise...
...base du montant du bilan apres reglement ; et

lorsqu'une transaction est reglee par le processus d'authentification, le montant du bilan apres reglement est calcule sur la base du bilan...
...authentifications sont compares dans chacune des transactions successives a regler par le biais d'un processus de non-authentification dans lequel la correspondance a l'authentification personnelle n'est pas requise...
...montant du reglement est etablie pour un reglement de transactions par le biais d'un processus de non-authentification dans lequel la correspondance a l'authentification personnelle n'est pas requise ; et

ledit moyen de comparaison et de determination determine qu'un processus illegal a ete execute avec la carte lorsqu'un montant de deboursement, qui va etre...
...du reglement.

6. Systeme selon l'une quelconque des revendications precedentes, dans lequel lorsque le processus d'authentification est invoque pour executer un versement ou pour regler une transaction, une somme d'argent est versee pour le processus d'authentification et est ecrite dans une zone predeterminee de la memoire de traitement des...
...autre dans la carte :

un moyen arithmetique destine a executer des calculs arithmetiques pour le processus d'authentification et un processus de non-authentification dans lequel la correspondance a l'authentification personnelle n'est pas requise...
...base d'un montant du bilan apres reglement lorsqu'une transaction est reglee par un processus d'authentification necessitant qu'une correspondance a l'authentification personnelle soit etablie, et le bilan...
...inferieur ou egal au bilan des authentifications memorise lorsque la transaction est reglee par le processus d'authentification,

mais seul le bilan des non-authentifications est mis a jour sur la base d'un montant du bilan apres reglement lorsque la transaction est reglee par un processus de non-authentification,

determiner qu'un processus illegal a ete execute avec la carte lorsqu'une comparaison du bilan des non-authentifications...
...11. Procede selon la revendication 10, dans lequel :

lorsqu'une transaction est reglee par un processus de non-authentification dans lequel la correspondance a l'authentification personnelle n'est pas requise... ..base du montant du bilan apres reglement ; et

lorsqu'une transaction est reglee par le processus d'authentification, le montant du bilan apres reglement est calcule sur la base du bilan... ..12. Procede selon la revendication 10, dans lequel :

lorsque la transaction est reglee par un processus de non-authentification dans lequel la correspondance a l'authentification personnelle n'est pas requise... ..base du montant du bilan apres reglement ; et

lorsqu'une transaction est reglee par le processus d'authentification, le montant du bilan apres reglement est calcule sur la base du bilan... ..authentifications sont compares dans chacune des transactions successives a regler par le biais d'un processus de non-authentification dans lequel la correspondance a l'authentification personnelle n'est pas requise... ..10, 11, 12 ou 13, comprenant en outre l'etape consistant a determiner qu'un processus illegal a ete execute avec la carte lorsqu'une limite de montant du reglement est... ..montant du reglement etant etablie pour le reglement de transaction par le biais d'un processus de non-authentification dans lequel la correspondance a l'authentification personnelle n'est pas requise.

15. Procede selon l'une quelconque des revendications 10 a 14, dans lequel le processus d'authentification est invoque pour executer un versement ou pour regler une transaction, une somme d'argent est versee pour le processus d'authentification et est ecrite sur une zone predeterminee de la memoire de traitement de... ..les etapes consistant :

executer a l'interieur de la carte des calculs arithmetiques pour le processus d'authentification et le processus de non-authentification dans lequel la correspondance a l'authentification personnelle n'est pas requise... ..montant du bilan apres reglement lorsqu'une transaction est reglee par le biais d'un processus d'authentification necessitant qu'une correspondance a une authentification personnelle soit etablie, et le bilan... ..egal au bilan des authentifications memorise lorsque la transaction est reglee par le biais du processus d'authentification ;

mais seul le bilan des non-authentifications est mis a jour sur la... ..montant du bilan apres reglement lorsque la transaction est reglee par le biais d'un processus de non-authentification,

determiner qu'un processus illegal a ete execute avec la carte lorsqu'une comparaison du bilan des non-authentifications... ..17 ou 18, dans lequel :

lorsqu'une transaction est reglee par le biais d'un processus de non-authentification dans lequel la correspondance a l'authentification personnelle n'est pas requise... ..base du montant du bilan apres reglement ; et

lorsqu'une transaction est reglee par le processus d'authentification, le montant du bilan apres reglement est calcule sur la base du bilan... ..selon la revendication 17 ou 18, dans lequel :

lorsqu'une transaction est reglee par un processus de non-authentification dans lequel la correspondance a l'authentification personnelle n'est pas requise... ..montant du bilan apres reglement ; et

lorsqu'une transaction est reglee par le biais du processus d'authentification, le montant du bilan apres reglement est calcule sur la base du bilan... ..authentifications sont compares dans chacune des transactions successives a regler par le biais d'un processus de non-authentification dans lequel la correspondance a l'authentification personnelle n'est pas requise... ..21, dans lequel ledit programme comprend en outre une procedure destinee a determiner qu'un processus illegal a ete execute avec la carte lorsqu'une limite de montant du reglement est... ..montant du reglement etant etablie pour un reglement de transactions par le biais d'un processus de non-authentification dans lequel la correspondance a l'authentification personnelle n'est pas requise... ..programme comprend en outre des procedures destinees a verser une somme d'argent pour le processus d'authentification et a ecrire la somme d'argent sur une zone predeterminee de la memoire de traitement de la non-authentification lorsque le processus d'authentification est invoque afin d'executer un versement ou de regler une transaction, la... ..etapes consistant a :

executer a l'interieur de la carte des calculs arithmetiques pour le processus d'authentification et un processus de non-authentification dans lequel la correspondance a l'authentification personnelle n'est pas requise... ...carte et une unite externe.

25. Systeme selon la revendication 1, comprenant en outre un terminal de transaction avec lequel il est possible de lire la carte.

26. Systeme selon la... ...systeme comprend en outre une unite arithmetique destinee a executer des calculs arithmetiques pour le processus d'authentification et le processus de non-authentification.

29. Systeme selon l'une quelconque des revendications 25 a 28, dans lequel ledit systeme comprend en outre une unite de synchronisation qui execute un processus de synchronisation apres qu'un versement ait ete execute par le processus d'authentification, le processus de synchronisation comprenant une etape consistant a mettre a jour un historique de deboursement par...

6/K/3 (Item 3 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Country	Number	Kind	Date
---------	--------	------	------

Abstract ...A2

Type	Pub. Date	Kind	Text
Available Text	Language	Update	Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: ...to (i) U.S. Application Serial No. 08/707,322 which is entitled "METHOD AND APPARATUS FOR CROSSTALK CANCELLATION" and (ii) U.S. Application Serial No.08/501,250 which is... ...the upstream and downstream signals occupy the same frequency bands and are separated by signal processing.

ANSI is producing another standard for subscriber line based transmission system, which is referred to... ...to as Fiber To The Curb (FTTC). The transmission medium from the "curb" to the customer is standard unshielded twisted-pair (UTP) telephone lines.

A number of modulation schemes have been... ...data signals are then supplied from the buffer 102 to a forward error correction (FEC) unit 104. The FEC unit 104 compensates for errors that are due to crosstalk noise, impulse noise, channel distortion, etc. The signals output by the FEC unit 104 are supplied to a data symbol encoder 106. The data symbol encoder 106 operates... ...encoded the data onto each of the frequency tones, an Inverse Fast Fourier Transform (IFFT) unit 112 modulates the frequency domain data supplied by the data symbol encoder 106 and produces... ...to digital signals. The digital signals are then supplied to a Fast Fourier Transform (FFT) unit 154 that demodulates the digital signals while converting the digital signals from a time domain... ...frequency domain. The demodulated digital signals are then supplied to a frequency domain equalizer (FEQ) unit 156. The FEQ unit 156 performs an equalization on the digital signals so the attenuation and phase are equalized... ...from each of the

frequency tones is then forwarded to the forward error correction (FEC) unit 164. The FEC unit 164 performs error correction of the data to produce corrected data. The corrected data is... a buffer 166. Thereafter, the data may be retrieved from the buffer 166 and further processed by the receiver 150. Alternatively, the received energy allocation table 160 could be supplied to and utilized by the FEQ unit 166.

The bit allocation tables and the energy allocation tables utilized in the conventional transmitter... components. For example, the transmitter 100 could add acyclic prefix to symbols after the IFFT unit 112, and the remote receiver 150 can then remove the cyclic prefix before the FFT unit 154. Also, the remote receiver 150 can provide a time domain equalizer (TEQ) unit between the ADC 152 and the FFT unit 154.

Most of the proposed VDSL/FTTC transmission schemes utilize frequency division duplexing (FDD) of... duplexing (TDD) of the upstream and downstream signals. More particularly, the time division duplexing is synchronized in this case such that periodic synchronized upstream and downstream communication periods do not overlap with one another. That is, the upstream and downstream communication periods for all of the wires that share a binder are synchronized. With this arrangement, all the very high speed transmissions within the same binder are synchronized and time division duplexed such that downstream communications are not transmitted at times that overlap... data is transmitted in either direction, separate the upstream and downstream communication periods. When the synchronized time division duplexed approach is used with DMT it is often referred to as synchronized DMT (SDMT).

A common feature of the above-mentioned transmission systems is that twisted-pair... As the speed of the data transmission increases, the problem worsens. The advantage of the synchronized TDD (such as SDMT) based data transmission is that crosstalk interference (NEXT interference) from other... format).

A data transmission system normally includes a central office and a plurality of remote units. Each remote unit communicates with the central office over a data link (i.e., channel) that is established between the central office and the particular remote unit. To establish such a data link, initialization processing is performed to initialize communications between the central office and each of the remote units. For purposes of the discussion to follow, a central office includes a central modem (or central unit) and a remote unit includes a remote modem. These modems are transceivers that facilitate data transmission between the central office and the remote unit. The central office thus normally includes a plurality of central side transceivers, each of which has a central side transmitter and a central side receiver, and the remote unit normally includes a remote side transceiver having a remote side transmitter and a remote side receiver.

One conventional frame synchronization technique required the transmission of a predetermined sequence of data which was received by a... a predetermined stored sequence of data to determine the adjustment required in order to yield synchronization. U.S. Patent No. 5,627,863 describes a frame synchronization approach suitable for systems (e.g., ADSL) using frequency division duplexing (FDD) or echo cancelling to provide duplexed operation. This frame synchronization technique requires a special start-up training sequence to obtain the frame synchronization. However, the described frame synchronization approach is not suitable for systems (e.g., synchronized TDD or SDMT) using time division duplexing because synchronization in time is not necessary for FDD or echo cancelling as it is with TDD... time-division duplexed (TDD) manner, the transmitters and receivers of the central office and remote units must be synchronized in time so that transmission and reception do not overlap in time. In a data... in the opposite direction occurs. Some transmission schemes divide upstream and downstream transmissions into smaller units referred to as frames. These frames may also be grouped into superframes that include a... during which it may transmit, and there are quiet periods (guard intervals) during which no unit must transmit. On channels subject to crosstalk (NEXT interference) between multiple connections, if time-division duplexing is used, synchronization must be established and maintained among all units so affected. An example is the VDSL service that uses the existing twisted pair telephone... modulation scheme. This scheme makes excellent use of time-division duplexing since a single FFT unit can be used during transmission and reception and avoids the need for two such FFT units, and other savings in the analog circuitry.

Conventional frame synchronization techniques are not only not well suited for synchronized TDD but also are unreliable when RF interference is present. Due to the potential for... ..or perhaps greater than, the desired receive signal power under some conditions. However, in a synchronized TDD system, it is important that synchronization be established and maintained so that crosstalk is mitigated and controlled and/or received data is accurately recovered.

Accordingly, there is a need for improved synchronization techniques for time-division duplexed systems.

SUMMARY OF THE INVENTION

Broadly speaking, the invention relates to improved techniques for synchronizing transmissions and receptions of a data transmission system utilizing time division duplexing. According to one aspect of the invention, the improved synchronization techniques utilize the time-varying nature of the energy of the received data to obtain synchronization. In one embodiment, the improved synchronization technique uses the output signals from a multicarrier modulation unit (e.g., FFT unit) and thus provides the ability to avoid frequency tones that are susceptible to RF interference. According to another aspect of the invention, the improved synchronization techniques utilize crosstalk interference levels to obtain synchronization. With the improved synchronization techniques, remote receivers in the data transmission system are able to synchronize to central transmitters, central receivers in the data transmission system are able to synchronize to remote transmitters, and central transmitters are able to synchronize with one another.

The invention can be implemented in numerous ways, including as an apparatus, system, method, or computer readable media. Several embodiments of the invention are discussed below.

As a method for adjusting... ..alignment error estimate using the edge detected in the plurality of consecutive frames. Additionally, the synchronization may thereafter be adjusted in accordance with the alignment error estimate. Optionally, the data transmission... ..set of the frames in the superframe transmit data in a second direction.

As a computer readable medium containing program instructions for adjusting an alignment for a first transceiver to receive... ..two-way data communication using time division duplexing, an embodiment of the invention includes: first computer readable code devices for measuring an energy amount for each of a plurality of consecutive frames of received data; second computer readable code devices for detecting an edge in the plurality of consecutive frames of the received data based on the measured energy amounts; and third computer readable code devices for computing an alignment error estimate using the edge detected in the... ..digital signals; an input buffer for temporarily storing the received digital signals; a multicarrier demodulation unit, the multicarrier demodulation Unit demodulates the received digital signals from the input buffer to frequency domain data for a plurality of different carrier frequencies; a frame synchronization unit, the frame synchronization unit synchronizes a receive frame boundary for the multicarrier demodulation unit based on the time-varying nature of the energy of the frequency domain data produced by the multicarrier demodulation unit; a bit allocation table, the allocation table stores bit allocation information used in transmitting data... ..for storing the decoded bits as recovered data. Preferably, the data transmission system is a

synchronized DMT system, and wherein the multicarrier demodulation unit includes a FFT unit.

As a receiver for a data transmission system using time division duplexing to alternate between... ..digital signals; an input buffer for temporarily storing the received digital signals; a multicarrier demodulation unit, the multicarrier demodulation unit demodulates the received digital signals from the input buffer to frequency domain data for a plurality of different carrier frequencies; frame synchronization means for synchronizing a receive frame boundary for the multicarrier demodulation unit based on the time-varying nature of the energy of the frequency domain data produced by the multicarrier demodulation unit; a bit allocation table, the allocation table stores bit allocation information used in transmitting data... ..plurality of transmitters at a central site where an external clock signal is

unavailable for synchronizing the transmitters, the transmitters transmit data in accordance with a superframe format including at least one quiet period, a method for synchronizing data transmissions by a given transmitter to other of the transmitters at the central site... ..at the central site; comparing the measured energy with a threshold amount; and modifying the synchronization for the transmissions by the given transmitter when the comparing indicates that the measured energy exceeds the threshold amount

As a computer readable medium containing program instructions for synchronizing data transmissions in a data transmission system having a plurality of transmitters at a central site where an external clock signal is unavailable for synchronizing the transmitters, the transmitters transmit data in accordance with a superframe format including at least one quiet period, an embodiment of the invention includes: first computer readable code devices for measuring the energy in the quiet period associated with a given transmitter due to data transmissions from other of the transmitters at the central site; second computer readable code devices for comparing the measured energy with threshold amount; and third computer readable code devices for modifying the synchronization for the transmissions by the given transmitter when the comparing indicates that the measured energy... ..amount.

The advantages of the invention are numerous. One advantage of the invention is that synchronization can be achieved even in the presence of radio frequency (RF) interference, such as due... ..that it is well suited for data transmission systems utilizing time division duplexing such as synchronized DMT or synchronized VDSL. Yet another advantage of the invention is that it is relatively insensitive to noise... ..telecommunications network suitable for implementing the invention;

FIG. 3 is a block diagram of a processing and distribution unit 300 according to an embodiment of the invention;

FIG. 4 is a diagram illustrating an... ..which a certain level of service is provided;

FIG. 5A is a flow-diagram of synchronization processing according to a basic embodiment of the invention;

FIG. 5B is a flow diagram of synchronization processing according to an embodiment of the invention;

FIGs. 6A and 6B are flow diagrams of synchronization processing according to a more detailed embodiment of the invention;

FIG. 7 is a flow diagram of edge detection processing according to an embodiment of the invention;

FIG. 8 is a flow diagram of alignment error estimation processing according to an embodiment of the invention;

FIGs. 9A and 9B represent diagrams of energy... ..to one embodiment of the invention; and

FIG. 12 is a flow diagram of a synchronization processing for synchronizing adjacent transmitters to compensate for small synchronization differences.

DETAILED DESCRIPTION OF THE INVENTION

The invention relates to improved techniques for synchronizing transmissions and receptions by a data transmission system utilizing time division duplexing. In one aspect of the invention, the improved synchronization techniques utilize the time-varying nature of the energy of the received data to obtain synchronization. In another aspect of the invention, the improved synchronization techniques utilize crosstalk interference levels to obtain synchronization. With the improved synchronization techniques, remote receivers in the data transmission system are able to synchronize to central transmitters, central receivers in the data transmission system are able to synchronize to remote transmitters, and central transmitters are able to synchronize with one another.

The synchronization required in a time-division duplex system requires that transmissions be synchronized with a superframe structure. Conventional time-domain methods which tend to correlate samples, such as... ..be of equal power to be desired signals. However, the invention provides accurate techniques to synchronize transmissions in a time-division duplex system even when RF interference renders the time domain signal unreliable. The frequency

domain approach to synchronization provided by the invention is able to obtain significant immunity from RF interference. In one embodiment, the improved synchronization technique preferably uses the output signals from a multicarrier modulation unit (FFT unit) and thus provides the ability to avoid frequency tones that are susceptible to radio frequency... posts to provide data transmission to and from the central office 202 to various remote units. In this exemplary embodiment, each of the distribution posts is a processing and distribution unit 204 (node). The processing and distribution unit 204 is coupled to the central office 202 by a high speed, multiplexed transmission line... fiber optic line. Typically, when the transmission line 206 is a fiber optic line, the processing and distribution unit 204 is referred to as an optical network unit (ONU). The central office 202 also usually interacts with and couples to other processing and distribution units (not shown) through high speed, multiplexed transmission lines 208 and 210, but only the operation of the processing and distribution unit 204 is discussed below. In one embodiment, the processing and distribution unit 204 includes one or more modems (central modems).

The processing and distribution unit 204 services a multiplicity of discrete subscriber lines 212-1 through 212-n. Each subscriber line 212 typically services a single end user. The end user has a remote unit suitable for communicating with the processing and distribution unit 204 at very high data rates. More particularly, a remote unit 214 of a first end user 216 is coupled to the processing and distribution unit 204 by the subscriber line 212-1, and a remote unit 218 of a second end user 220 is coupled to the processing and distribution unit 204 by the subscriber line 212-n. The remote units 214 and 218 include a data communications system capable of transmitting data to and receiving data from the processing and distribution unit 204. In one embodiment, the data communication systems are modems. The remote units 214 and 218 can be incorporated within a variety of different devices, including for example, a telephone, a television, a monitor, a computer, a conferencing unit, etc. Although FIG. 2 illustrates only a single remote unit coupled to a respective subscriber line, it should be recognized that a plurality of remote units can be coupled to a single subscriber line. Moreover, although FIG. 2 illustrates the processing and distribution unit 204 as being centralized processing, it should be recognized that the processing need not be centralized and could be performed independently for each of the subscriber lines 212.

The subscriber lines 212 serviced by the processing and distribution unit 204 are bundled in a shielded binder 222 as the subscriber lines 212 leave the processing and distribution unit 204. The shielding provided by the shielded binder 222 generally serves as a good insulator... shielded binder 222 and is coupled directly or indirectly to the end user's remote units. The "drop" portion of the subscriber line between the respective remote unit and the shielded binder 222 is normally an unshielded, twisted-pair wire. In most applications... transmitted are allocated. The telecommunications network 200, for example, is particularly well suited for a synchronized TDD transmission system (e.g., synchronized VDSL or SDMT) offering different levels of service.

Hence, referring to the SDMT transmission system... 2, data transmissions over all lines 212 in the shielded binder 222 associated with the processing and distribution unit 204 need to be synchronized. As such, all active lines emanating from the processing and distribution unit 204 could be transmitting in the same direction (i.e., downstream or upstream) so as to substantially eliminate NEXT interference.

FIG. 3 is a block diagram of a processing and distribution unit 300 according to an embodiment of the invention. For example, the data processing and distribution unit 300 is a detailed implementation of the processing and distribution unit 204 illustrated in FIG. 2.

The data processing and distribution unit 300 includes a processing unit 302 that receives data and sends data over a data link 304. The data link... coupled to a fiber optic cable of a telephone network or a cable network. The processing unit 302 needs to operate to synchronize the various processed transmissions and receptions of the processing unit 302. The data processing and distribution unit 300 further includes a bus arrangement 308 and a plurality of analog cards 310. The output of the processing unit 302 is coupled to the bus arrangement 308. The bus arrangement 308 together with the processing unit 302 thus direct output data from the processing unit 302 to the appropriate analog cards 310 as well as direct input from the analog cards 310 to the processing

unit 302. The analog cards 310 provide analog circuitry utilized by the processing and distribution unit 300 that is typically more efficiently performed with analog components than using digital processing by the processing unit 302. For example, the analog circuitry can include filters, transformers, analog-to-digital converters or... to the fifty (50) lines. In one embodiment, the lines are twisted-pair wires. The processing Unit 302 may be a general-purpose computing device such as a digital signal processor (DSP) or a dedicated special purpose device. The bus arrangement 308 may take many arrangements... be a single card or circuitry that supports multiple lines.

In a case where the processing is not centralized, the processing unit 302 in FIG. 3 can be replaced by modems for each of the lines. The processing for each of the lines can then be performed independently for each of the lines... circuitry.

The NEXT interference problem occurs on the lines proximate to the output of the processing and distribution unit 300. With respect to the block diagram illustrated in FIG. 3, the NEXT interference is... power differential (between transmitted and received signals). In other words, from the output of the processing and distribution unit 300 the lines travel towards the remote units. Usually, most of the distance is within a shielded binder that would, for example, hold... frames are inserted between the upstream and the downstream bursts to allow the channel to settle before the direction of transmission is changed.

FIG. 4 is a diagram illustrating an exemplary... 16 symbols downstream; 1 quiet period; 2 symbols upstream; and 1 quiet period.

With proper synchronization at a central unit (processing and distribution unit 204 or processing unit 302) and uniform superframe formats, synchronized transmissions of equal duration are provided for all lines within a binder. Accordingly, the NEXT interference problem is effectively eliminated. The synchronization of the central unit and the remote units is also important for accurate data recovery. These synchronizations are needed in synchronized VDSL and SDMT systems. According to the invention, improved synchronization techniques are described below with respect to FIGs. 5-12.

FIG. 5A is a flow diagram of synchronization processing 500 according to a basic embodiment of the invention. Initially, the synchronization processing 500 measures 502 energy in n consecutive frames of received data. An alignment error estimate... based on the measured energy values for the n consecutive frames. Following block 504, the synchronization processing 500 is complete and ends.

FIG. 5B is a flow diagram of synchronization processing 550 according to an embodiment of the invention. Initially, the synchronization processing 550 measures 552 energy in n consecutive frames of received data. Next, an edge is... then computed 556 from the position of the edge that has been detected. Thereafter, the synchronization processing 550 is able to adjust 558 its synchronization reference in accordance with the alignment error estimate. Following block 558, the synchronization processing 550 is complete and ends.

By determining and adjusting synchronization of receivers of the remote units to transmissions from a central unit in accordance with the synchronization processing 500 or 550, the remote units are able to establish synchronization with the central unit. Once synchronized the central unit and the remote units are able to share a channel (transmission line) in a time-division duplexed manner. Also, the synchronization processing 500 or 550 is suitable for determining and adjusting synchronization of receivers at the central unit with transmissions from the remote units.

FIGs. 6A and 6B are flow diagrams of synchronization processing 600 according to a more detailed embodiment of the invention. Once the synchronization processing 600 is initiated, FFT outputs are obtained 602 for n consecutive frames of received data... forward the received data to an analog-to-digital converter and then to a FFT unit, such as illustrated in FIG. 1B. Hence, the FFT outputs may be obtained from the output of the FFT unit. The FFT outputs are frequency domain signals.

Next, the FFT outputs that are susceptible to RF interference are dropped 604. The remaining FFT outputs are then used for subsequent processing. Typically, a frame includes a plurality of different frequency tones. Each of the frequency tones... ..to the RF interference due to amateur radio users. In the case of a remote unit of a synchronized multicarrier VDSL system, where a frame has 256 frequency tones, frequency tones 6 through 40... ..attenuated because lower frequency tones have less attenuation, and therefore sufficient to obtain a reliable synchronization result. Hence, in one embodiment, frequency tones 6 through 40 from each of n consecutive frames are used for subsequent processing.

Next, energy values for the n consecutive frames of the remaining FFT outputs are determined... ..frequency tones 6 through 40 are being utilized, then the corresponding outputs from the FFT unit are obtained and then converted to energy values and summed together so as to produce...frame can be obtained by summing the squared moduli of all outputs of the FFT unit that are in use. Alternatively, the energy values could be obtained by summing the energies... ..interference.

Once the energy values for the n consecutive frames have been determined 606, the synchronization processing 600 detects 608 a burst edge within the received data based on the determined energy... ..the beginning (or end) of a received transmission from the transmitter and additionally identifies a synchronization for the frame. A trailing edge within the received data and/or characteristics of the... ..boundary setting. In particular, from the determined energy values in the burst edge, the remote unit synchronization processing 600 is able to determine the alignment error for a frame (i.e., error in frame synchronization). Typically, the alignment error is estimated as a fraction of a frame. Thereafter, the frame... ..be adjusted 612 in accordance with the alignment error estimate.

Once adjusted 612, the frame synchronization should be established. However, preferably, the synchronization processing 600 continues to confirm that the synchronization has been achieved. Specifically, following block 612, a decision block 614 determines whether the absolute... ..threshold. If the alignment error estimate is not less than a predetermined threshold, then the synchronization processing 600 returns to repeat block 602 and subsequent blocks so as to iteratively reduce the... ..received transmission and/or the number of frames in the burst. Following block 616, the synchronization processing 600 is complete and ends.

Normally, when the frame synchronization is adjusted 612 by a significant amount, the alignment error estimate is greater than the predetermined threshold. Hence, the synchronization processing 600 will repeat and should produce a small alignment error amount which is less than the predetermined threshold. Then, the synchronization processing 600 is able to proceed to block 616. Alternatively, the decision block 614 can be... ..with a high degree of confidence.

FIG. 7 is a flow diagram of edge detection processing 700 according to an embodiment of the invention. The edge detection processing 700 describes additional details on the block 608 in FIG. 6A where the burst edge is detected. The edge detection processing 700 initially computes 702 successive energy differences for the n determined energy values. These successive... .. $j+1$) are then stored 706 for later retrieval. Following block 706, the edge detection processing 700 is complete and the processing returns to block 610 of the synchronization processing 600.

FIG. 8 is a flow diagram of alignment error estimation processing 800 according to an embodiment of the invention. The alignment error estimation processing 800 describes additional details on the block 610 in FIG. 6A where an alignment error estimate is determined. The alignment error estimation processing 800 initially determines 802 a difference amount from the energy values at indices $(j+1...$..this embodiment, the alignment error estimate represents a fractional part of a frame. Accordingly, the synchronization of the receiver to the data transmission unit would be off by this fractional part of the frame. Following block 804, the alignment error estimation processing 800 is complete and the processing returns to block 612 of the synchronization processing 600.

FIGs. 9A and 9B represent diagrams of energy values (e) and energy difference values... ..in FIG. 7.

As seen in FIGs. 9A and 9B, the receiver is not properly synchronized with the incoming transmitted data from a remotely located transmitter. In particular, the beginning of... burst of data received from the transmitter begins somewhere within frame 6. To be properly synchronized, the burst of data from the transmitter would begin exactly at the beginning of frame... an alignment adjustment has been made in accordance with the invention, that is, with proper synchronization. In FIG. 10A a diagram 1000 indicates a burst of data between frames 6 and... inferred (9 frames), and the superframe format can be identified (9-1-9-1).

During synchronization the successive differences in the energy values observed in each frame of the superframe will... the end of a burst. According to one embodiment of the invention, the edge detection processing adjusts the frame alignment so that the maximum difference is increased, the right-hand neighboring energy difference is forced to zero. When synchronization has been obtained, the result is as shown in FIG. 10B. Note that the edge detection processing is relatively insensitive to the absolute amplitudes being observed. The successive differences approach requires only... the cyclic prefix because the removal of the cyclic prefix drops samples useful for frame

synchronization but thus unavailable from the FFT unit. One technique to resolve this dead; zone where the frame has 512 samples and the... estimates to get a combined energy estimate which is then used in the burst detection processing.

The synchronization processing discussed above is generally applicable to remote side and central side synchronization. For synchronization processing at a remote unit, a receiver at the remote unit acquires and maintains synchronization with data transmissions (bursts) with a transmitter of the central unit. As for synchronization processing at a central unit, a receiver at the central unit acquires and maintains synchronization with data transmissions (bursts) with a transmitter of a remote unit. In one embodiment, the synchronization is managed by setting or adjusting receive frame alignment for the recovery of data transmissions... of a line (or channel), the time at which an upstream transmission from a remote unit reaches a central unit will vary and will appear to be late by the length of the round-trip delay if no correction is made. Accordingly, the central unit needs to adjust its receive frame alignment so that the correct receive samples are used in the receiver at the central unit. The processing carried out at the central unit to adjust its receive frame alignment is similar to the synchronization processing discussed above for the remote unit. Generally, the energy in upstream frames being received is measured over a number of frames corresponding to the length of the upstream transmission burst from the remote unit. These energy values are used to identify the start of the upstream transmission burst and... align the receive frame boundary pointer with the frames of data received from the remote unit.

FIG. 11 is a block diagram of a receiver 1100 according to one embodiment of... may be used in either or both of the central office transceiver and the remote unit transceiver.

The receiver 1100 receives analog signals 1102 that have been transmitted over a channel... then supplied to an input buffer 1106 that temporarily stores the digital signals. The FFT unit 1108 retrieves a frame of data from the input buffer 1106 in accordance with a... pointer 1110, and then produces frequency domain signals.

In accordance with the invention, the FFT unit 1108 outputs the frequency domain signals 1112 to a frame synchronization unit 1114. The frame synchronization unit 1114 operates to perform the synchronization processing discussed above with reference to FIGs. 5-10B. The frame synchronization unit 1114 outputs an alignment error estimate 1116 to a controller 1118. The controller 1118 then adjusts the receive frame boundary pointer 1110 for accessing the received data from the input buffer 1106. Hence, the frame synchronization unit 1114 provides for frame synchronization in the time domain duplexed transmission system in a manner that is substantially immune from RF... receiver 1100. The controller 1118, for example, controls the receiver 1100 to perform the initialization processing and to monitor steady-state data transmission. For example, the controller 1118 can be implemented by a digital signal processor, a microprocessor or microcontroller, or specialized circuitry. In the case where the receiver 1100 forms... a plurality of transceivers, or individually provided for each transmitter

and receiver. Likewise, the frame synchronization Unit 1114 can be implemented by a digital signal processor, a microprocessor or microcontroller, or specialized circuitry.

Returning to the receive data path the frequency domain signals 1112 output by the FFT unit 1108 are then equalized by the FEQ unit 1120. The equalized signals are then Supplied to a data symbol decoder 1122. The data... bit and energy allocation table 1124. The decoded data is then supplied to the FEC unit 1126 and then stored in an output buffer 1128. Thereafter, recovered data 1130 (stored decoded... For example, when a corresponding transmitter adds a cyclic prefix to symbols after an IFFT unit, the receiver 1100 can remove the cyclic prefix before the FFT unit 1108. Also, the receiver 1100 can provide a time domain equalizer (TEQ) unit between the ADC 1104 and the FFT unit 1106. Additional details on TEQ units are contained in U.S. Patent No. 5,285,474 and U.S. Application Serial... TIME DOMAIN EQUALIZATION, which are hereby incorporated by reference.

Moreover, the invention provides techniques to synchronize transmissions at a central side (i.e., central unit). With synchronized transmissions at the central side, the NEXT interference is substantially eliminated, provided all lines of... if the transmissions from the central side over lines in a binder are not properly synchronized, the NEXT interference is a substantial impediment to efficient and accurate operation of the data... central side transmissions. If the NEXT interference is not strong enough to be detected for synchronization purposes, then it will be assumed to be insignificant during reception, and therefore synchronization is not necessary.

Conventionally, the various transmitters at the central side can synchronize with one another by all using a common master clock supplied to the central side... positioned at slightly greater positions from the master clock source so as to cause small synchronization differences between the various transmissions. Hence, the synchronization techniques according to the invention can also be used to synchronize various transmissions at the central side.

FIG. 12 is a flow diagram of a synchronization processing 1200 for synchronizing adjacent transmitters to compensate for small synchronization differences. If these small synchronization differences were to go uncorrected, over time the degree of the lack of synchronization worsens. The synchronization processing 1200 initially measures 1202 energy received from other transmitters at the central side. Here, during... interference is detected, it is known that the transmitters at the central side are not synchronized. Hence, the timing alignment at the transmitter is modified 1202 in order to synchronize its alignment with respect to the other transmitters at the central side. For example, the... Following block 1206 or following block 1204 when the predetermined threshold is not exceeded, the synchronization processing 1200 is complete and ends.

The synchronization processing 1200 is performed by all the transceivers at the central side. By repeating the synchronization processing 1200, gradually the alignment will reach a more less steady state, particularly if adjustments to... illustrated in FIG. 4, the superframe format has two quiet periods 404 and 408. The synchronization processing 1200 uses one of the two quiet periods 404 and 408. When a receiver at... transceiver adjust its timing alignment at the central side, it may inform the corresponding remote unit of the change so that it also modifies its timing alignment. This notification to the remote can, for example, be performed over an overhead channel.

The synchronization technique needs to distinguish downstream NEXT interference from upstream FEXT interference. This can be achieved... distinguishing feature is detected (greater than some threshold) it means that the clock in this unit is running faster than the interfering transmitter's clock.

The adjustment to the synchronization can be to modify the clock frequency of the particular transceiver's clock, such as... then insertion of 1 sample per superframe (11,040 samples) will be sufficient to monitor synchronization. If the central side transceivers can only insert, the central side transceivers will reach a... transmission.

The advantages of the invention are numerous. One advantage of the invention is that synchronization can be achieved even in the presence of radio frequency (RF) interference, such as due... that it is well suited for data

transmission systems utilizing time division duplexing such as synchronized DMT or synchronized VDSL. Yet another advantage of the invention is that it is relatively insensitive to background... Also included is a method as recited hereinabove, wherein the first transceiver is a remote unit and the second transceiver is a central unit.

Also included is a method as recited hereinabove, wherein the second transceiver is a remote unit and the first transceiver is a central unit.

Also included is a method as recited hereinabove, wherein the energy amount is a power... of the frames of the superframe structure including the cyclic prefix.

Also included is a computer readable medium containing program instructions for adjusting an alignment for a first transceiver to receive... with a data transmission system providing two-way data communication using time division duplexing, the computer readable medium comprising: first computer readable code devices for measuring an energy amount for each of a plurality of consecutive frames of received data; and second computer readable code devices for computing an alignment error estimate based on the measured energy amounts.

Also included is the computer readable medium as recited hereinabove, wherein the second computer readable medium comprises: computer readable code devices for detecting an edge in the plurality of consecutive frames of the received data based on the measured energy amounts; and computer readable code devices for determining the alignment error estimate using the edge detected in the plurality of consecutive frames.

Also included is the computer readable medium as recited hereinabove, wherein the second computer readable medium further comprises: computer readable code devices for computing successive energy differences in the plurality of the measured energy amounts; and computer readable code devices for identifying a largest one of the successive energy differences, the largest one of the successive energy difference corresponding to a burst edge.

Also included is the computer readable medium as recited hereinabove, wherein the second computer readable medium comprises: computer readable code devices for identifying a prior energy difference and a subsequent energy difference, the... one of the successive differences immediately following the largest one of the successive energy differences; computer readable code devices for determining the alignment error estimate based on the prior energy difference and the subsequent energy difference.

Also included is the computer readable medium as recited hereinabove, wherein the data transmission system transmits data using a superframe... of the frames contain a cyclic prefix for the superframe structure, and wherein the first computer readable code devices for measuring of the energy amounts comprises: computer readable code for measuring energy amounts of a first set of consecutive frames of received data for the superframe structure; computer readable code for measuring energy amounts of a second set of consecutive frames of received... consecutive frames being offset from and overlapped with the first set of the consecutive frames; and computer readable code for combining together the energy amounts from respective consecutive frames from the first and second sets of the consecutive frames to produce the energy amounts for the second computer readable code devices.

Also included is the computer readable medium as recited hereinabove, wherein the combining determines mean energy amounts for each of the frames of the superframe structure including the cyclic prefix.

Also included is the computer readable medium as recited hereinabove, wherein the number of frames in the first and second... digital signals; an input buffer for temporarily storing the received digital signals; a multicarrier demodulation unit, the multicarrier demodulation unit demodulates the received digital signals from the input buffer to frequency domain data for a plurality of different carrier frequencies; a frame synchronization unit, the frame synchronization unit synchronizes a receive frame boundary for the multicarrier demodulation unit based on the time-varying nature of the energy of the frequency domain data produced by the multicarrier demodulation unit; a

bit allocation table, the allocation table stores bit allocation information used in transmitting data...bits as recovered data.

Also included is the receiver as recited hereinabove, wherein the frame synchronization unit determines an alignment adjustment amount, wherein the receiver further comprises a controller for controlling overall operation of the receiver, the controller receives the alignment adjustment amount from the frame synchronization unit and accordingly adjusts a receive frame boundary pointer for the input buffer.

Also included is the receiver as recited hereinabove, wherein at least one of the frame synchronization unit and the controller are implemented by a processor.

Also included is the receiver as recited hereinabove, wherein the frame synchronization unit is a processor.

Also included is the receiver as recited hereinabove, wherein the received data signals undesirably include radio frequency interference, and wherein the frame synchronization unit ignores the portion of the frequency domain data that overlaps with frequency ranges of the... ..Also included is the receiver as recited hereinabove, wherein the data transmission system is a synchronized DMT system, and wherein the multicarrier demodulation unit includes a FFT unit.

Also included is a receiver for a data transmission system using time division duplexing to... ..digital signals; an input buffer for temporarily storing the received digital signals; a multicarrier demodulation unit, the multicarrier demodulation unit demodulates the received digital signals from the input buffer to frequency domain data for a plurality of different carrier frequencies; frame synchronization means for synchronizing a frame boundary for the multicarrier demodulation unit based on the time-varying nature of the energy of the frequency domain data produced by the multicarrier demodulation unit; a bit allocation table, the allocation table stores bit allocation information used in transmitting data... ..in accordance with a superframe format including at least one quiet period, a method for synchronizing data transmissions by a given transmitter to other of the transmitters at the central site... ..central site; (b) comparing the measured energy with a threshold amount; and (c) modifying the synchronization for the transmissions by the given transmitter when the comparing (b) indicates that the measured... ..is a multicarrier data transmission system, and wherein an external clock signal is unavailable for synchronizing the transmitters, and wherein the modifying (c) comprises adjusting timing alignment to reduce crosstalk interference... ..of the transmitters and not due to the incoming data receptions.

Also included is a computer readable medium containing program instructions for synchronizing data transmissions in a data transmission systems having a plurality of transmitters at a central site where an external clock signal is unavailable for synchronizing the transmitters, the transmitters transmit data in accordance with a superframe format including at least one quiet period, the computer readable medium comprising: first computer readable code devices for measuring the energy in the quiet period associated with a given transmitter due to data transmissions from other of the transmitters at the central site; second computer readable code devices for comparing the measured energy with a threshold amount; and third computer readable code devices for modifying the synchronization for the transmissions by the given transmitter when the comparing indicates that the measured energy exceeds the threshold amount.

Also included is the computer readable medium as recited hereinabove, wherein the data transmission system is a multicarrier data transmission... ..and the transmitters are part of transceivers at the central site, and wherein the third computer readable code devices operates to adjust timing alignment to reduce crosstalk interference.

The many features...

Specification: ...the upstream and downstream signals occupy the same frequency bands and are separated by signal processing.

ANSI is producing another standard for subscriber line based transmission system, which is referred to... ..to as Fiber To The Curb (FTTC). The transmission medium from the "curb" to the customer is standard unshielded twisted-pair (UTP) telephone lines.

A number of modulation schemes have been... ..data signals are then supplied from the buffer 102 to a forward error correction (FEC) unit 104. The FEC unit 104 compensates for errors that are due to crosstalk noise, impulse noise, channel distortion, etc. The signals output by the FEC unit 104 are supplied to a data symbol encoder 106. The data symbol encoder 106 operates... ..encoded the data onto each of the frequency tones, an Inverse Fast Fourier Transform (IFFT) unit 112 modulates the frequency domain data supplied by the data symbol encoder 106 and produces... ..to digital signals. The digital signals are then supplied to a Fast Fourier Transform (FFT) unit 154 that demodulates the digital signals while converting the digital signals from a time domain... ..frequency domain. The demodulated digital signals are then supplied to a frequency domain equalizer (FEQ) unit 156. The FEQ unit 156 performs an equalization on the digital signals so the attenuation and phase are equalized... ..from each of the frequency tones is then forwarded to the forward error correction (FEC) unit 164. The FEC unit 164 performs error correction of the data to produce corrected data. The corrected data is... ..a buffer 166. Thereafter, the data may be retrieved from the buffer 166 and further processed by the receiver 150. Alternatively, the received energy allocation table 160 could be supplied to and utilized by the FEQ unit 166.

The bit allocation tables and the energy allocation tables utilized in the conventional transmitter... ..For example, the transmitter 100 could add a cyclic prefix to symbols after the IFFT unit 112, and the remote receiver 150 can then remove the cyclic prefix before the FFT unit 154. Also, the remote receiver 150 can provide a time domain equalizer (TEQ) unit between the ADC 152 and the FFT unit 154.

Most of the proposed VDSL/FTTC transmission schemes utilize frequency division duplexing (FDD) of... ..duplexing (TDD) of the upstream and downstream signals. More particularly, the time division duplexing is synchronized in this case such that periodic synchronized upstream and downstream communication periods do not overlap with one another. That is, the upstream and downstream communication periods for all of the wires that share a binder are synchronized. With this arrangement, all the very high speed transmissions within the same binder are synchronized and time division duplexed such that downstream communications are not transmitted at times that overlap... ..data is transmitted in either direction, separate the upstream and downstream communication periods. When the synchronized time division duplexed approach is used with DMT it is often referred to as synchronized DMT (SDMT).

A common feature of the above-mentioned transmission systems is that twisted-pair... ..As the speed of the data transmission increases, the problem worsens. The advantage of the synchronized TDD (such as SDMT) based data transmission is that crosstalk interference (NEXT interference) from other... ..format).

A data transmission system normally includes a central office and a plurality of remote units. Each remote unit communicates with the central office over a data link (i.e., channel) that is established between the central office and the particular remote

unit. To establish such a data link, initialization processing is performed to initialize communications between the central office and each of the remote units. For purposes of the discussion to follow, a central office includes a central modem (or central unit) and a remote unit includes a remote modem. These modems are transceivers that facilitate data transmission between the central office and the remote unit. The central office thus normally includes a plurality of central side transceivers, each of which has a central side transmitter and a central side receiver, and the remote unit normally includes a remote side transceiver having a remote side transmitter and a remote side receiver.

One conventional frame synchronization technique required the transmission of a predetermined sequence of data which was received by a... ..a predetermined stored sequence of data to determine the adjustment required in order

to yield synchronization. U.S. Patent No. 5,627,863 describes a frame synchronization approach suitable for systems (e.g., ADSL) using frequency division duplexing (FDD) or echo cancelling to provide duplexed operation. This frame synchronization technique requires a special start-up training sequence to obtain the frame synchronization. However, the described frame synchronization approach is not suitable for systems (e.g., synchronized TDD or SDMT) using time division duplexing because synchronization in time is not necessary for FDD or echo cancelling as it is with TDD... ..time-division duplexed (TDD) manner, the transmitters and receivers of the central office and remote units must be synchronized in time so that transmission and reception do not overlap in time. In a data... ..in the opposite direction occurs. Some transmission schemes divide upstream and downstream transmissions into smaller units referred to as frames. These frames may also be grouped into superframes that include a... ..during which it may transmit, and there are quiet periods (guard intervals) during which no unit must transmit. On channels subject to crosstalk (NEXT interference) between multiple connections, if time-division duplexing is used, synchronization must be established and maintained among all units so affected. An example is the VDSL service that uses the existing twisted pair telephone... ..modulation scheme. This scheme makes excellent use of time-division duplexing since a single FFT unit can be used during transmission and reception and avoids the need for two such FFT units, and other savings in the analog circuitry.

A prior art synchronized TDD subscriber Line system is disclosed by WO 97/03506 mentioned supra.

Conventional frame synchronization techniques are not only not well suited for synchronized TDD but also are unreliable when RF interference is present. Due to the potential for... ..or perhaps greater than, the desired receive signal power under some conditions. However, in a synchronized TDD system, it is important that synchronization be established and maintained so that crosstalk is mitigated and controlled and/or received data is accurately recovered.

Accordingly, there is a need for improved synchronization techniques for time-division duplexed systems.

SUMMARY OF THE INVENTION

Broadly speaking, the invention relates to improved techniques for synchronizing transmissions and receptions of a data transmission system utilizing time division duplexing. According to one aspect of the invention, the improved synchronization techniques utilize the time-varying nature of the energy of the received data to obtain synchronization. In one embodiment, the improved synchronization technique uses the output signals from a multicarrier modulation unit (e.g., FFT unit) and thus provides the ability to avoid frequency tones that are susceptible to RF interference. According to another aspect of the invention, the improved synchronization techniques utilize crosstalk interference levels to obtain synchronization. With the improved synchronization techniques, remote receivers in the data transmission system are able to synchronize to central transmitters, central receivers in the data transmission system are able to synchronize to remote transmitters, and central transmitters are able to synchronize with one another.

The invention can be implemented in numerous ways, including as an apparatus, system, method or computer readable media. Several embodiments of the invention are discussed below.

There is provided a method... ..the alignment error estimate is less than the threshold amount.

There is also provided a computer readable medium containing program instructions for adjusting an alignment for a first transceiver to receive... ..digital signals; an input buffer for temporarily storing the received digital signals; a multicarrier demodulation unit, said multicarrier demodulation unit demodulates the received digital signals from said input buffer to frequency domain data for a plurality of different carrier frequencies; a frame synchronization unit, said frame synchronization unit synchronizes a receive frame boundary for said multicarrier demodulation unit by measuring an energy amount for each of a plurality of consecutive frames of received... ..1128) for storing the decoded bits as recovered data.

There is provided a method for synchronizing data transmissions by a given transmitter to other of the transmitters at a central site... (a) comparing the measured energy with a threshold amount; (b) comparing the measured energy with a threshold amount; and (c) modifying the synchronization for the transmissions by the given transmitter when said comparing (b) indicates that the measured energy exceeds the threshold amount.

There is also provided a computer readable medium containing program instructions for synchronizing data transmissions in a data transmission systems having a plurality of transmitters at a central site where an external clock signal is unavailable for synchronizing the transmitters, the transmitters transmit data in accordance with a superframe format including at least... (a) at the central site; comparing the measured energy with a threshold amount; and modifying the synchronization for the transmissions by the given transmitter when said comparing indicates that the measured energy... exceeds the threshold amount.

The advantages of the invention are numerous. One advantage of the invention is that synchronization can be achieved even in the presence of radio frequency (RF) interference, such as due... (a) that it is well suited for data transmission systems utilizing time division duplexing such as synchronized DMT or synchronized VDSL. Yet another advantage of the invention is that it is relatively insensitive to noise... telecommunications network suitable for implementing the invention;

FIG. 3 is a block diagram of a processing and distribution unit 300 according to an embodiment of the invention;

FIG. 4 is a diagram illustrating an... (a) which a certain level of service is provided;

FIG. 5A is a flow diagram of synchronization processing according to a basic embodiment of the invention;

FIG. 5B is a flow diagram of synchronization processing according to an embodiment of the invention;

FIGs. 6A and 6B are flow diagrams of synchronization processing according to a more detailed embodiment of the invention;

FIG. 7 is a flow diagram of edge detection processing according to an embodiment of the invention;

FIG. 8 is a flow diagram of alignment error estimation processing according to an embodiment of the invention;

FIGs. 9A and 9B represent diagrams of energy... (a) to one embodiment of the invention; and

FIG. 12 is a flow diagram of a synchronization processing for synchronizing adjacent transmitters to compensate for small synchronization differences.

DETAILED DESCRIPTION OF THE INVENTION

The invention relates to improved techniques for synchronizing transmissions and receptions by a data transmission system utilizing time division duplexing. In one aspect of the invention, the improved synchronization techniques utilize the time-varying nature of the energy of the received data to obtain synchronization. In another aspect of the invention, the improved synchronization techniques utilize crosstalk interference levels to obtain synchronization. With the improved synchronization techniques, remote receivers in the data transmission system are able to synchronize to central transmitters, central receivers in the data transmission system are able to synchronize to remote transmitters, and central transmitters are able to synchronize with one another.

The synchronization required in a time-division duplex system requires that transmissions be synchronized with a superframe structure. Conventional time-domain methods which tend to correlate samples, such as... (a) be of equal power to be desired signals. However, the invention provides accurate techniques to synchronize transmissions in a time-division duplex system even when RF interference renders the time domain signal unreliable. The frequency domain approach to synchronization provided by the invention is able to obtain significant immunity from RF interference. In one embodiment, the improved synchronization technique preferably uses the output signals from a multicarrier modulation unit (FFT unit) and thus provides the ability to avoid frequency tones that are susceptible to

radio frequency... posts to provide data transmission to and from the central office 202 to various remote units. In this exemplary embodiment, each of the distribution posts is a processing and distribution unit 204 (node). The processing and distribution unit 204 is coupled to the central office 202 by a high speed, multiplexed transmission line... fiber optic line. Typically, when the transmission line 206 is a fiber optic line, the processing and distribution unit 204 is referred to as an optical network unit (ONU). The central office 202 also usually interacts with and couples to other

processing and distribution units (not shown) through high speed, multiplexed transmission lines 208 and 210, but only the operation of the processing and distribution unit 204 is discussed below. In one embodiment, the processing and distribution unit 204 includes one or more modems (central modems).

The processing and distribution unit 204 services a multiplicity of discrete subscriber lines 212-1 through 212-n. Each subscriber line 212 typically services a single end user. The end user has a remote unit suitable for communicating with the processing and distribution unit 204 at very high data rates. More particularly, a remote unit 214 of a first end user 216 is coupled to the processing and distribution unit 204 by the subscriber line 212-1, and a remote unit 218 of a second end user 220 is coupled to the processing and distribution unit 204 by the subscriber line 212-n. The remote units 214 and 218 include a data communications system capable of transmitting data to and receiving data from the processing and distribution unit 204. In one embodiment, the data communication systems are modems. The remote units 214 and 218 can be incorporated within a variety of different devices, including for example, a telephone, a television, a monitor, a computer, a conferencing unit, etc. Although FIG. 2 illustrates only a single remote unit coupled to a respective subscriber line, it should be recognized that a plurality of remote units can be coupled to a single subscriber line. Moreover, although FIG. 2 illustrates the processing and distribution unit 204 as being centralized processing, it should be recognized that the processing need not be centralized and could be performed independently for each of the subscriber lines 212.

The subscriber lines 212 serviced by the processing and distribution unit 204 are bundled in a shielded binder 222 as the subscriber lines 212 leave the processing and distribution unit 204. The shielding provided by the shielded binder 222 generally serves as a good insulator... shielded binder 222 and is coupled directly or indirectly to the end user's remote units. The "drop" portion of the subscriber line between the respective remote unit and the shielded binder 222 is normally an unshielded, twisted-pair wire. In most applications... transmitted are allocated. The telecommunications network 200, for example, is particularly well suited for a synchronized TDD transmission system (e.g., synchronized VDSL or SDMT) offering different levels of service.

Hence, referring to the SDMT transmission system... 2, data transmissions over all lines 212 in the shielded binder 222 associated with the processing and distribution unit 204 need to be synchronized. As such, all active lines emanating from the processing and distribution unit 204 could be transmitting in the same direction (i.e., downstream or upstream) so as to substantially eliminate NEXT interference.

FIG. 3 is a block diagram of a processing and distribution unit 300 according to an embodiment of the invention. For example, the data processing and distribution unit 300 is a detailed implementation of the processing and distribution unit 204 illustrated in FIG. 2.

The data processing and distribution unit 300 includes a processing unit 302 that receives data and sends data over a data link 304. The data link... coupled to a fiber optic cable of a telephone network or a cable network. The processing unit 302 needs to operate to synchronize the various processed transmissions and receptions of the processing unit 302. The data processing and distribution unit 300 further includes a bus arrangement 308 and a plurality of analog cards 310. The output of the processing unit 302 is coupled to the bus arrangement 308. The bus arrangement 308 together with the processing unit 302 thus direct output data from the processing unit 302 to the appropriate analog cards 310 as well as direct input from the analog cards 310 to the processing unit 302. The analog cards 310 provide analog circuitry utilized by the processing and distribution unit 300 that is typically more

efficiently performed with analog components than using digital processing by the processing unit 302. For example, the analog circuitry can include filters, transformers, analog-to-digital converters or... to the fifty (50) lines. In one embodiment, the lines are twisted-pair wires. The processing unit 302 may be a general-purpose computing device such as a digital signal processor (DSP) or a dedicated special purpose device. The bus arrangement 308 may take many arrangements... be a single card or circuitry that supports multiple lines.

In a case where the processing is not centralized, the processing unit 302 in FIG. 3 can be replaced by modems for each of the lines. The processing for each of the lines can then be performed independently for each of the lines... circuitry.

The NEXT interference problem occurs on the lines proximate to the output of the processing and distribution unit 300. With respect to the block diagram illustrated in FIG. 3, the NEXT interference is... power differential (between transmitted and received signals). In other words, from the output of the processing and distribution unit 300 the lines travel towards the remote units. Usually, most of the distance is within a shielded binder that would, for example, hold... frames are inserted between the upstream and the downstream bursts to allow the channel to settle before the direction of transmission is changed.

FIG. 4 is a diagram illustrating an exemplary... 16 symbols downstream; 1 quiet period; 2 symbols upstream; and 1 quiet period.

With proper synchronization at a central unit (processing and distribution unit 204 or processing unit 302) and uniform superframe formats, synchronized transmissions of equal duration are provided for all lines within a binder. Accordingly, the NEXT interference problem is effectively eliminated. The synchronization of the central unit and the remote units is also important for accurate data recovery. These synchronizations are needed in synchronized VDSL and SDMT systems. According to the invention, improved synchronization techniques are described below with respect to FIGs. 5-12.

FIG. 5A is a flow diagram of synchronization processing 500 according to a basic embodiment of the invention. Initially, the synchronization processing 500 measures 502 energy in n consecutive frames of received data. An alignment error estimate... based on the measured energy values for the n consecutive frames. Following block 504, the synchronization processing 500 is complete and ends.

FIG. 5B is a flow diagram of synchronization processing 550 according to an embodiment of the invention. Initially, the synchronization processing 550 measures 552 energy in n consecutive frames of received data. Next, an edge is... then computed 556 from the position of the edge that has been detected. Thereafter, the synchronization processing 550 is able to adjust 558 its synchronization reference in accordance with the alignment error estimate. Following block 558, the synchronization processing 550 is complete and ends.

By determining and adjusting synchronization of receivers of the remote units to transmissions from a central unit in accordance with the synchronization processing 500 or 550, the remote units are able to establish synchronization with the central unit. Once synchronized the central unit and the remote units are able to share a channel (transmission line) in a time-division duplexed manner. Also, the synchronization processing 500 or 550 is suitable for determining and adjusting synchronization of receivers at the central unit with transmissions from the remote units.

FIGs. 6A and 6B are flow diagrams of synchronization processing 600 according to a more detailed embodiment of the invention. Once the synchronization processing 600 is initiated, FFT outputs are obtained 602 for n consecutive frames of received data... forward the received data to an analog-to-digital converter and then to a FFT unit, such as illustrated in FIG. 1B. Hence, the FFT outputs may be obtained from the output of the FFT unit. The FFT outputs are frequency domain signals.

Next, the FFT outputs that are susceptible to RF interference are dropped 604. The remaining FFT outputs are then used for subsequent processing. Typically, a frame includes a plurality of different frequency tones. Each of the

frequency tones... ..to the RF interference due to amateur radio users. In the case of a remote unit of a synchronized multicarrier VDSL system, where a frame has 256 frequency tones, frequency tones 6 through 40... ..attenuated because lower frequency tones have less attenuation, and therefore sufficient to obtain a reliable synchronization result. Hence, in one embodiment, frequency tones 6 through 40 from each of n consecutive frames are used for subsequent processing.

Next, energy values for the n consecutive frames of the remaining FFT outputs are determined... ..frequency tones 6 through 40 are being utilized, then the corresponding outputs from the FFT unit are obtained and then converted to energy values and summed together so as to produce... ..frame can be obtained by summing the squared moduli of all outputs of the FFT unit that are in use. Alternatively, the energy values could be obtained by summing the energies... ..interference.

Once the energy values for the n consecutive frames have been determined 606, the synchronization processing 600 detects 608 a burst edge within the received data based on the determined energy... ..the beginning (or end) of a received transmission from the transmitter and additionally identifies a synchronization for the frame. A trailing edge within the received data and/or characteristics of the... ..boundary setting. In particular, from the determined energy values in the burst edge, the remote unit synchronization processing

600 is able to determine the alignment error for a frame (i.e., error in frame synchronization). Typically, the alignment error is estimated as a fraction of a frame. Thereafter, the frame... ..be adjusted 612 in accordance with the alignment error estimate.

Once adjusted 612, the frame synchronization should be established. However, preferably, the synchronization processing 600 continues to confirm that the synchronization has been achieved. Specifically, following block 612, a decision block 614 determines whether the absolute... ..threshold. If the alignment error estimate is not less than a predetermined threshold, then the synchronization processing 600 returns to repeat block 602 and subsequent blocks so as to iteratively reduce the... ..received transmission and/or the number of frames in the burst. Following block 616, the synchronization processing 600 is complete and ends.

Normally, when the frame synchronization is adjusted 612 by a significant amount, the alignment error estimate is greater than the predetermined threshold. Hence, the synchronization processing 600 will repeat and should produce a small alignment error amount which is less than the predetermined threshold. Then, the synchronization processing 600 is able to proceed to block 616. Alternatively, the decision block 614 can be... ..with a high degree of confidence.

FIG. 7 is a flow diagram of edge detection processing 700 according to an embodiment of the invention. The edge detection processing 700 describes additional details on the block 608 in FIG. 6A where the burst edge is detected. The edge detection processing 700 initially computes 702 successive energy differences for the n determined energy values. These successive... .. $j+1$) are then stored 706 for later retrieval. Following block 706, the edge detection processing 700 is complete and the processing returns to block 610 of the synchronization processing 600.

FIG. 8 is a flow diagram of alignment error estimation processing 800 according to an embodiment of the invention. The alignment error estimation processing 800 describes additional details on the block 610 in FIG. 6A where an alignment error estimate is determined. The alignment error estimation processing 800 initially determines 802 a difference amount from the energy values at indices $(j+1$this embodiment, the alignment error estimate represents a fractional part of a frame. Accordingly, the synchronization of the receiver to the data transmission unit would be off by this fractional part of the frame. Following block 804, the alignment error estimation processing 800 is complete and the processing returns to block 612 of the synchronization processing 600.

FIGs. 9A and 9B represent diagrams of energy values (e) and energy difference values... ..in FIG. 7.

As seen in FIGs. 9A and 9B, the receiver is not properly synchronized with the incoming transmitted data from a remotely located transmitter. In particular, the beginning of... burst of data received from the transmitter begins somewhere within frame 6. To be properly synchronized, the burst of data from the transmitter would begin exactly at the beginning of frame... an alignment adjustment has been made in accordance with the invention, that is, with proper synchronization. In FIG. 10A a diagram 1000 indicates a burst of data between frames 6 and... inferred (9 frames), and the superframe format can be identified (9-1-9-1).

During synchronization the successive differences in the energy values observed in each frame of the superframe will... the end of a burst. According to one embodiment of the invention, the edge detection processing adjusts the frame alignment so that the maximum difference is increased, the right-hand neighboring energy difference is forced to zero. When synchronization has been obtained, the result is as shown in FIG. 10B. Note that the edge detection processing is relatively insensitive to the absolute amplitudes being observed. The successive differences approach requires only... the cyclic prefix because the removal of the cyclic prefix drops samples useful for frame synchronization but thus unavailable from the FFT unit. One technique to resolve this dead-zone where the frame has 512 samples and the... estimates to get a combined energy estimate which is then used in the burst detection processing.

The synchronization processing discussed above is generally applicable to remote side and central side synchronization. For synchronization processing at a remote unit, a receiver at the remote unit acquires and maintains synchronization with data transmissions (bursts) with a transmitter of the central unit. As for synchronization processing at a central unit, a receiver at the central unit acquires and maintains synchronization with data transmissions (bursts) with a transmitter of a remote unit. In one embodiment, the synchronization is managed by setting or adjusting receive frame alignment for the recovery of data transmissions... of a line (or channel), the time at which an upstream transmission from a remote unit reaches a central unit will vary and will appear to be late by the length of the round-trip delay if no correction is made. Accordingly, the central unit needs to adjust its receive frame alignment so that the correct receive samples are used in the receiver at the central unit. The processing carried out at the central unit to adjust its receive frame alignment is similar to the synchronization processing discussed above for the remote unit. Generally, the energy in upstream frames being received is measured over a number of frames corresponding to the length of the upstream transmission burst from the remote unit. These energy values are used to identify the start of the upstream transmission burst and... align the receive frame boundary pointer with the frames of data received from the remote unit.

FIG. 11 is a block diagram of a receiver 1100 according to one embodiment of... may be used in either or both of the central office transceiver and the remote unit transceiver.

The receiver 1100 receives analog signals 1102 that have been transmitted over a channel... then supplied to an input buffer 1106 that temporarily stores the digital signals. The FFT unit 1108 retrieves a frame of data from the input buffer 1106 in accordance with a... pointer 1110, and then produces frequency domain signals.

In accordance with the invention, the FFT unit 1108 outputs the frequency domain signals 1112 to a frame synchronization unit 1114. The frame synchronization unit 1114 operates to perform the synchronization processing discussed above with reference to FIGs. 5-10B. The frame synchronization unit 1114 outputs an alignment error estimate 1116 to a controller 1118. The controller 1118 then adjusts the receive frame boundary pointer 1110 for accessing the received data from the input buffer 1106. Hence, the frame synchronization unit 1114 provides for frame synchronization in the time domain duplexed transmission system in a manner that is substantially immune from... receiver 1100. The controller 1118, for example, controls the receiver 1100 to perform the initialization processing and to monitor steady-state data transmission. For example, the controller 1118 can be implemented by a digital signal processor, a microprocessor or microcontroller, or specialized circuitry. In the case where the receiver 1100 forms... a plurality of transceivers, or individually provided for each transmitter and receiver. Likewise, the frame synchronization unit 1114 can be implemented by a digital signal processor, a microprocessor or microcontroller, or specialized circuitry.

Returning to the receive data path the frequency domain signals 1112 output by the FFT unit 1108 are then equalized by the FEQ unit 1120. The equalized signals are then supplied to a data symbol decoder 1122. The data... bit and energy allocation table 1124. The decoded data is then supplied to the FEC unit 1126 and then stored in an output buffer 1128. Thereafter, recovered data 1130 (stored decoded... For example, when a corresponding transmitter adds a cyclic prefix to symbols after an IFFT unit, the receiver 1100 can remove the cyclic prefix before the FFT unit 1108. Also, the receiver 1100 can provide a time domain equalizer (TEQ) unit between the ADC 1104 and the FFT unit 1106. Additional details on TEQ units are contained in U.S. Patent No. 5,285,474.

Moreover, the invention provides techniques to synchronize transmissions at a central side (i.e., central unit). With synchronized transmissions at the central side, the NEXT interference is substantially eliminated, provided all lines of... if the transmissions from the central side over lines in a binder are not properly synchronized, the NEXT interference is a substantial impediment to efficient and accurate operation of the data... central side transmissions. If the NEXT interference is not strong enough to be detected for synchronization purposes, then it will be assumed to be insignificant during reception, and therefore synchronization is not necessary.

Conventionally, the various transmitters at the central side can synchronize with one another by all using a common master clock supplied to the central side... positioned at slightly greater positions from the master clock source so as to cause small synchronization differences between the various transmissions. Hence, the synchronization techniques according to the invention can also be used to synchronize various transmissions at the central side.

FIG. 12 is a flow diagram of a synchronization processing 1200 for synchronizing adjacent transmitters to compensate for small synchronization differences. If these small synchronization differences were to go uncorrected, over time the degree of the lack of synchronization worsens. The synchronization processing 1200 initially measures 1202 energy received from other transmitters at the central side. Here, during... interference is detected, it is known that the transmitters at the central side are not

synchronized. Hence, the timing alignment at the transmitter is modified 1202 in order to synchronize its alignment with respect to the other transmitters at the central side. For example, the... Following block 1206 or following block 1204 when the predetermined threshold is not exceeded, the synchronization processing 1200 is complete and ends.

The synchronization processing 1200 is performed by all the transceivers at the central side. By repeating the synchronization processing 1200, gradually the alignment will reach a more less steady state, particularly if adjustments to... illustrated in FIG. 4, the superframe format has two quiet periods 404 and 408. The synchronization processing 1200 uses one of the two quiet periods 404 and 408. When a receiver at... transceiver adjust its timing alignment at the central side, it may inform the corresponding remote unit of the change so that it also modifies its timing alignment. This notification to the remote can, for example, be performed over an overhead channel.

The synchronization technique needs to distinguish downstream NEXT interference from upstream FEXT interference. This can be achieved... distinguishing feature is detected (greater than some threshold) it means that the clock in this unit is running faster than the interfering transmitter's clock.

The adjustment to the synchronization can be to modify the clock frequency of the particular transceiver's clock, such as... then insertion of 1 sample per superframe (11,040 samples) will be sufficient to monitor synchronization. If the central side transceivers can only insert, the central side transceivers will reach a... transmission.

The advantages of the invention are numerous. One advantage of the invention is that synchronization can be achieved even in the presence of radio frequency (RF) interference, such as due... that it is well suited for data

transmission systems utilizing time division duplexing such as synchronized DMT or synchronized VDSL. Yet another advantage of the invention is that it is relatively insensitive to background...

Claims: ...d) indicates that the alignment error estimate is less than the threshold amount.

7. A computer readable medium containing program instructions for adjusting an alignment for a first transceiver to receive... with a data transmission system providing two-way data communication using time division duplexing, said computer readable medium comprising:

first computer readable code devices for measuring an energy amount for each of a plurality of consecutive frames of received data; and

second computer readable code devices for computing an alignment error estimate based on the measured energy amounts... digital signals;

an input buffer for temporarily storing the received digital signals;

a multicarrier demodulation unit, said multicarrier demodulation unit demodulates the received digital signals from said input buffer to frequency domain data for a plurality of different carrier frequencies;

a frame synchronization unit, said frame synchronization unit synchronizes a receive frame boundary for said multicarrier demodulation unit based on the time-varying nature of the energy of the frequency domain data produced by said multicarrier demodulation unit;

a bit allocation table, said allocation table stores bit allocation information used in transmitting data... in accordance with a superframe format including at least one quiet period, a method for synchronizing data transmissions by a given transmitter to other of the transmitters at the central site... central site;

(b) comparing the measured energy with a threshold amount; and

(c) modifying the synchronization for the transmissions by the given transmitter when said comparing (b) indicates that the measured energy exceeds the threshold amount.

10. A computer readable medium containing program instructions for synchronizing data transmissions in a data transmission systems having a plurality of transmitters at a central site where an external clock signal is unavailable for synchronizing the transmitters, the transmitters transmit data in accordance with a superframe format including at least one quiet period, said computer readable medium comprising:

first computer readable code devices for measuring the energy in the quiet period associated with a given transmitter due to data transmissions from other of the transmitters at the central site;

second computer readable code devices for comparing the measured energy with a threshold amount; and

third computer readable code devices for modifying the synchronization for the transmissions by the given transmitter when said comparing indicates that the measured energy...

Claims: ...d) indicates that the alignment error estimate is less than the threshold amount.

7. A computer readable medium containing program instructions for adjusting a time alignment for a first transceiver to... signals;

an input buffer (1106) for temporarily storing the received digital signals;

a multicarrier demodulation unit (1108), said multicarrier demodulation unit being adapted to demodulate the received digital signals from said input buffer to frequency domain data for a plurality of different carrier frequencies;

a frame synchronization unit (1114), said frame synchronization unit being adapted to synchronize a receive frame boundary for said multicarrier demodulation unit by measuring an energy amount for each of a plurality of consecutive frames of received... ..output buffer (1128) for storing the decoded bits as recovered data.

9. A method for synchronizing data transmissions by a given transmitter to other of the transmitters at a central site... ..central site;

(b) comparing the measured energy with a threshold amount; and

(c) modifying the synchronization for the transmissions by the given transmitter when said comparing (b) indicates that the measured energy exceeds the threshold amount.

10. A computer readable medium containing program instructions for synchronizing data transmissions in a data transmission systems having a plurality of transmitters at a central site where an external clock signal is unavailable for synchronizing the transmitters, the transmitters transmit data in accordance with a superframe format including at least... ..at the central site;

comparing the measured energy with a threshold amount; and

modifying the synchronization for the transmissions by the given transmitter when said comparing indicates that the measured energy...

6/K/4 (Item 4 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Start code detecting apparatus for video data stream

Country	Number	Kind	Date
---------	--------	------	------

Abstract ...A2

Type	Pub. Date	Kind	Text
Available Text	Language	Update	Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: INTRODUCTION

The present invention is directed to improvements in methods and apparatus for decompression which operates to decompress and/or decode a plurality of differently encoded input... ..of the well known standards known as JPEG, MPEG and H.261.

A serial pipeline processing system of the present invention comprises a single two-wire bus used for carrying... ..to a plurality of adaptive decompression circuits and the like positioned as a reconfigurable pipeline processor.

PRIOR ART

One prior art system is described in United States Patent No. 5,216,724. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and a host processor. The device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus.

United States Patent No. 4,785... ...a known quad tree data structure.

United States Patent No. 5,122,875 discloses an apparatus for encoding/decoding an HDTV signal. The apparatus includes a compression circuit responsive to high definition video source signals for providing hierarchically layered... ...compressed video data of relatively greater and lesser importance to image reproduction respectively. A transport processor, responsive to the high and low priority codeword sequences, forms high and low priority transport... ...United States Patent No. 5,168,356 discloses a video signal encoding system that includes apparatus for segmenting encoded video data into transport blocks for signal ...in respective transport blocks.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening. This is accomplished by an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ...frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ...oscillators and the receiver can continuously receive each channel, then the receiver need not be synchronized with the transmitter. An FFT algorithm implements a fast discrete approximation to the continuous case in which the receiver synchronizes to the first frame and then acquires subsequent frames every frame period. The frame period... ...increasing the amount of data transmitted.

United States Patent No. 5,212,742 discloses an apparatus and method for processing video data for compression/decompression in real-time. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and host processor. Lastly, the device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus. The method handles assigning portions of the image for each of the processors to operate upon.

United States Patent No. 5,231,484 discloses a system and method... ...MPEG standards. Included are three cooperating components or subsystems that operate to variously adaptively pre-process the incoming digital motion video sequences, allocate bits to the pictures in a sequence, and... ...States Patent No. 5,267,334 discloses a method of removing frame redundancy in a computer system for a sequence of moving images. The method comprises detecting a first scene change... ...facing" keyframe or intraframe, and it is normally present in CCITT compressed video data. The process then comprises generating at least one intermediate compressed frame, the at least one intermediate compressed frame containing difference information from the first image for at least one image following...change, known as a "backward-facing" keyframe. The first keyframe and the at least one intermediate compressed frame are linked for forward play, and the second keyframe and the intermediate compressed frames are linked in reverse for reverse play. The intraframe may also be used... ...of complete scene information.

United States Patent No. 5,276,513 discloses a first circuit apparatus, comprising a given number of prior-art image-pyramid stages, together with a second circuit apparatus, comprising the same given number of novel motion-vector stages, perform cost-effective hierarchical motion analysis (HMA) in real-time, with minimum system processing delay and/or employing minimum system processing delay and/or employing minimum hardware structure. Specifically, the first and second circuit apparatus, in response to relatively high-resolution image data from an ongoing input series of successive... ...a relatively high frame rate (e.g., 30 frames per second), derives, after

a certain processing-system delay, an ongoing output series of successive given pixel-density vector-data frames that... ..of successive image frames.

United States Patent No. 5,283,646 discloses a method and apparatus for enabling a real-time video encoding system to accurately deliver the desired number of... ..desired bit allocations.

The article, Chong, Yong M., A Data-Flow Architecture for Digital Image Processing, Wescon Technical Papers: No. 2 Oct./Nov. 1984, discloses a real-time signal processing system specifically designed for image processing. More particularly, a token based data-flow architecture is disclosed wherein the tokens are of... ..width having a fixed width address field. The system contains a plurality of identical flow processors connected in a ring fashion. The tokens contain a data field, a control field and a tag. The tag field of the token is further broken down into a processor address field and an identifier field. The processor address field is used to direct the tokens to the correct data-flow processor, and the identifier field is used to label the data such that the data-flow processor knows what to do with the data. In this way, the identifier field acts as an instruction for the data-flow processor. The system directs each token to a specific data-flow processor using a module number (MN). If the MN matches the MN of the particular stage... ..to locate the decoder in the preceding stage in order to pre-decode complex decoding processing and to alleviate critical path problems in the logic circuit. The elastic nature of the...of block signal in most cases.

United States Patent No. 4,903,018 discloses a process and data processing system for compressing and expanding structurally associated multiple data sequences. The process is particular to data sets in which an analysis is made of the structure in... ..data series on the basis of the order number of these data elements. The data processing system for performing the processes includes a storage matrix (26) and an index storage (28) having line addresses of the... ..the final actual video.

United States Patent No. 5,060,242 discloses an image signal processing system DPCM encodes the signal, then Huffman and run length encodes the signal to produce... ..tightly packed without gaps for efficient transmission without loss of any data. The tightly packed apparatus has a barrel shifter with its shift modulus controlled by an accumulator receiving code word... ..OR gate is connected to the shifter, while a register is connected to the gate. Apparatus for processing a tightly packed and decorrelated digital signal has a barrel shifter and accumulator for unpacking... ..an inverse DCPM decoder.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening is accomplished by use of an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ..frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ..kernel matrix.

United States Patent No. 5,231,486 discloses a high definition video system processes a bitstream including high and low priority variable length coded Data words. The coded Data... ..packed High Priority Data and packed Low Priority Data by means of respective data packing units. The coded Data is continuously applied to both packing units. High Priority and Low Priority Length words indicating the bit lengths of high priority and... ..States Patent No. 5,287,178 discloses a video signal encoding system includes a signal processor for segmenting encoded video data into transport blocks having a header section and a packed data section. The system also includes reset control apparatus for releasing resets of system components, after a global system reset, in a prescribed non-simultaneous phased sequence to enable signal processing to commence in the prescribed sequence. The phased reset release sequence begins when valid data...United States Patent No. 5,142,380 to Sakagami et al. discloses an image compression apparatus suitable for use with still images such as those formed by electronic still cameras using... ..and Q.

United States Patent No. 5,193,002 to Guichard et al. disclosed an apparatus for coding/decoding image signals in real time in conjunction with the CCITT standard H.261. A digital signal processor carries out direct quantization and reverse quantization.

United States Patent No. 5,241,383 to Chen et al. describes an apparatus with a pseudo-constant bit rate video coding achieved by an adjustable quantization parameter. The... relates to an improved pipeline system having an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along the pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the processing stages in the pipeline and interacting with selected stages in the pipeline for control data and/or combined control-data functions among the processing stages, so that the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing. In accordance with the invention, the processing stages may be configurable in response to recognition of at least one token. One of the processing stages may be a Start Code Detector which receives the input and generates and/or... and resetting the system, and a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards. The present invention... data and having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video pictures of varying data size can be controlled. Also in accordance with the invention, a processing stage receives the input data stream, the stage including means for recognizing specified bit stream patterns, whereby the processing stage facilitates random access and error recovery. The invention may also include a means for... invention also includes an inverse modeller stage, an inverse discrete cosine transform stage, and a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for processing data.

In addition, the present invention relates to an improved pipeline system having a Huffman...pipeline stage that incorporates a two-wire transfer control and also shows two consecutive pipeline processing stages with the two-wire transfer control;

Figures. 5a and 5b taken together depict one... shown in Figures. 8a and 8b.

Figure 10 is a block diagram of a reconfigurable processing stage;

Figure 11 is a block diagram of a spatial decoder;

Figure 12 is a... decoder including the prediction filters;

Figure 18 is a pictorial representation of the prediction filtering process;

Figure 19 shows a generalized representation of the macroblock structure;

Figure 20 shows a generalized... buffer;

Figure 25 is a pictorial diagram illustrating prediction data offset from the block being processed;

Figure 26 is a pictorial diagram illustrating prediction data offset by (1,1);

Figure 27...in general terms, the present invention provides an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along a pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the stages in the pipeline and interacting with selected stages in the pipeline for control, data and/or combined control-data functions among the processing stages, whereby the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing.

Each of the processing stages in the pipeline may include both primary and secondary storage, and the stages in processing stages for performance of functions or position independent of the processing stages for performance of functions.

In a pipeline machine, in accordance with the invention, the... ..altered by interfacing with the stages, and the tokens may interact with all of the processing stages in the pipeline or only with some but less than all of said processing stages. The tokens in the pipeline may interact with adjacent processing stages or with non-adjacent processing stages, and the tokens may reconfigure the processing stages. Such tokens may be position dependent for some functions and position independent for other... ..be Huffman coded.

In the improved pipeline machine, the tokens may be generated by a processing stage. Such pipeline tokens may include data for transfer to the processing stages or the tokens may be devoid of data. Some of the tokens may be identified as DATA tokens and provide data to the processing stages in the pipeline, while other tokens are identified as control tokens and only condition the processing stages in the pipeline, such conditioning including reconfiguring of the processing stages. Still other tokens may provide both data and conditioning to the processing stages in the pipeline. Some of said tokens may identify coding standards to the processing stages in the pipeline, whereas other tokens may operate independent of any coding standard among the processing stages. The tokens may be capable of successive alteration by the processing stages in the pipeline.

In accordance with the invention, the interactive flexibility of the tokens in cooperation with the processing stages facilitates greater functional diversity of the processing stages for resident structure in the pipeline, and the flexibility of the tokens facilitates system... ..or alteration. The tokens may be capable of facilitating a plurality of functions within any processing stage in the pipeline. Such pipeline tokens may be either hardware based or software based... ..system bandwidth in the pipeline. The tokens may provide data and control simultaneously to the processing stages in the pipeline.

The invention may include a pipeline processing machine for handling plurality of separately encoded bit streams arranged as a single serial bit... ..and for passing unrecognized control tokens along the pipeline, and a reconfigurable decode and parser processing means responsive to a recognized control token for reconfiguring a particular stage to handle an... ..be a pipeline system and the Start Code Detector may be positioned as the first processing stage in the pipeline.

The present invention also provides, in a system having a plurality of processing stages, a universal adaptation unit in the form of an interactive interfacing token for control and/or data functions among the processing stages, the token being a PICTURE(underscore)START code token for indicating that the start... ..The token may also be a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards.

The CODING(underscore... ..standard as JPEG, and/or any other appropriate picture standard. At least some of the processing stages reconfigure in response to the CODING(underscore)STANDARD token.

One of the processing stages in the system may be a Huffman decoder and parser and, upon receipt of... ..Data stage, and the parser stage may send an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... ..video data, having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video... ..controlled.

The system may include a spatial decoder having a two-wire interface interconnecting processing stages, the interface enabling serial processing for data and parallel processing for control.

As previously indicated, the system may further include a ROM having separate stored... ..of a plurality of picture standards, the programs being selectable by a token to facilitate processing for a plurality of different picture standards.

The spatial decoder system also includes a token... ..decoding stage and a parser stage for sending an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... ..The present invention also provides a pipeline system having an input data stream, and a processing stage for receiving the input data stream, the stage including means for recognizing specified bit... ..whereby said stage facilitates random access and error recovery. In accordance with the invention, the processing stage may be a start code detector and the bit stream patterns may include start... ..token and padding insures uniformity of word size. In accordance with the invention, a reconfigurable processing stage may be provided as a spatial decoder and the padding means adds to picture... ..that if the DATA token has less than the predetermined length, the padder circuit adds units of data to the DATA token until the predetermined length is achieved. A bypass circuit...1 tokens into a buffer, having a second predetermined width.

The invention also provides an apparatus for providing a time delay to a group of compressed pictures, the pictures corresponding to... ..and capable of delaying the words of data, is in communication with a control circuit intermediate the counter circuit and the inverse modeller circuit, the control circuit also communicating with the... ..inverse modeller stage and an inverse discrete cosine transform stage, the improvement characterized by a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for processing data.

In accordance with the invention, the token may be a QUANT(underscore)TABLE token for causing the

processing stage to generate a quantization table.

The present invention also provides a Huffman decoder for... ..of bits used to represent an item of data.

DECODER: An embodiment of a decoding process.

DECODING (PROCESS): The process defined in this specification that reads an input coded bitstream and produces decoded pictures or... ..the same order in which they were presented at the input of the encoder.

ENCODING (PROCESS): A process, not specified in this specification, that reads a stream of input pictures or audio samples...to provide an estimate of the pel value or data element currently being decoded.

RECONFIGURABLE PROCESS STAGE (RPS): A stage, which in response to a recognized token, reconfigures itself to perform various operations.

SLICE: A series of macroblocks.

TOKEN: A universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions... ..indicates that the corresponding stage holds valid data, i.e., data that is to be processed in one of the pipeline stages. After processing (which may involve nothing more than a simple transfer without manipulation of the data) valid... ..present invention may be used with any number of pipeline stages. Furthermore, data may be processed in more than one stage and the processing time for different stages can differ.

In addition to clock and data signals (described below... ..other system. For example, the last pipeline stage may pass its data on to subsequent processing circuitry. The ACCEPT signal, which is illustrated as the lower of the two lines connecting... ..the minimum disturbance possible to other pipeline stages. Succeeding pipeline stages are allowed to continue processing and, therefore, this means that gaps open up in the stream of data following the... ..The data in the pipeline is encoded such that many different types of data are processed in the pipeline. This encoding accommodates data packets of variable size and the size of...the other hand, it may generate itself, all or part of the data to be processed in the pipeline. Indeed, as is explained below, a "stage" may contain arbitrary processing circuitry, including none at all (for simple passing of data) or entire systems (for example... ..values zero and 255 may not be used.

If such a picture were to be processed in a pipeline built in the practice of the present invention, then one of these... data must not be written over since it is data that must be saved for processing or use in a downstream device e.g., a pipeline stage, a device or a... connected to the pipeline upstream contains data D4 that is to be transferred into and processed in the pipeline. ...pipeline, in accordance with the preferred embodiments of the present invention, to "fill up" empty processing stages is highly advantageous since the processing stages in the pipeline thereby become decouple from one another. In other words, even though... data can be transferred into the pipeline and between stages even when one or more processing stages is blocked.

In the embodiment shown in Fig. 1, it is assumed that the...propagate all the way back to the beginning of the pipeline if there is some intermediate stage that is able to accept new data.

In the embodiment illustrated in Fig. 1...has been mentioned. It is to be further understood that each pipeline stage may also process the data it has received arbitrarily before passing it between its internal storage elements or... the portion of the pipeline that contains input and output storage elements and that arbitrarily processes data stored in its storage elements.

Furthermore, the "device" downstream from ...valid data, but also when a stage requires more than one clock phase to finish processing its data. This also can occur when it creates valid data in one or both... control the passage of data between adjacent storage elements. The VALID signal may also be processed in an analogous manner.

A great advantage of the two-wire interface (one wire for... In addition, two extra latches and a small number of gates are preferably added to process the ACCEPT and VALID signals that are associated with the data latches in each half... application so requires. The interface in accordance with this embodiment can also be used to process analog signals.

As discussed previously, while other conventional timing arrangements may be used, the interface... circuit B1, which may be provided to convert output data from input latch LDIN into intermediate data, which is then later loaded in an output data latch LDOUT, which comprises the... is connected either directly as an input to the validation output latch LVOUT, or via intermediate logic devices or circuits that may alter the signal.

Similarly, the output validation signal QVOUT... to the input of the validation input latch QVIN of the following stage, or via intermediate devices or logic circuits, which may alter the validation signal. This output QVIN ...word.

Preferred Data Structure - "tokens"

In the sample application shown in Fig. 4, each stage processes all input data, since there is no control circuitry that excludes any stage from allowing... are connected together in a relatively simple configuration. The simplest configuration is a pipeline of processing steps. For example, in the one shown in Fig. 1. The use of tokens, however...flows from left to right in the diagram. Data enters the machine and passes into processing Stage A. This may or may not modify the data and it then passes the... advantage of the tokens is their ability to achieve this kind of communication. Since any processing stage that does not recognize a token simply passes it on unaltered to the next... is transmitted along with the address and data fields in each token so that a processing stage can pass on a token (which can be of arbitrary length) without having to... be the first word of a new token.

Note that although the simple pipeline of processing stages is particularly useful, it will be appreciated that tokens may be applied to more complicated configurations of processing elements. An example of a more complicated processing element is described below.

It is not necessary, in accordance with the present invention, to... has extension bits. An example of this is a token that activates a stage that processes video quantization values stored in a quantization table (typically a memory device). For example, a...turn, is of great importance in video data pipeline systems since it ensures that all processing stages can be continuously running at full bandwidth.

In accordance to the present invention, in... some other chips in the set. This is advantageous both from the perspective of a customer and from that of a chip manufacturer. Even if modifications mean that all chips are... the end of a token (and hence the start of the next token) to be processed correctly (including simple non-manipulative transfer), even if the token is not recognized by the... block diagram of a pipeline stage whose function is as follows. If the stage is processing a predetermined token (known in this example as the DATA token), then it will duplicate... the address field of the DATA token. If, on the other hand, the stage is processing any other kind of token, it will delete every word. The overall effect is that... respective output signals:

In the duplication stage, the output from the data latch LDIN forms intermediate data referred to as MID(underscore)DATA. This intermediate data word is loaded into the data output latch LDOUT only when an intermediate acceptance signal (labeled "MID(underscore)ACCEPT" in Fig. 8a) is set HIGH.

The portion of... data. These include a "DATA(underscore)TOKEN" signal that indicates that the circuitry is currently processing a valid DATA Token, and a NOT(underscore)DUPLICATE signal which is used to control duplication of data. When the circuitry is processing a DATA Token, the NOT(underscore)DUPLICATE signal toggles between a HIGH and a LOW... the token to be duplicated once (but no more times). When the circuitry is not processing a valid DATA Token then the NOT(underscore)DUPLICATE signal is held in a HIGH state. Accordingly, this means that the token words that are being processed are not duplicated.

As Fig. 8a illustrates, the upper six bits of 8-bit intermediate data word and the output signal QI1 from the latch LI1 form inputs to a... explained further below.

Latch LO1 performs the function of latching the last value of the intermediate extension bit (labeled "MID(underscore)EXTN" and as signal S4), and it loads this value... and the DATA(underscore)TOKEN signal will become "0", indicating that the circuitry is not processing a DATA token.

If QI1 is "0" and SO is "0", thereby indicating a DATA... phase and the DATA(underscore)TOKEN signal will become "1", indicating that the circuitry is processing a DATA token.

The NOT(underscore)DUPLICATE signal (the output signal QO3) is similarly loaded... LVOUT at the same time that MID(underscore)DATA is loaded into LDOUT and the intermediate extension bit (signal S4) is loaded into LEOUT. Signal S5 is also combined with the... above. This has the effect that all tokens except the one that causes the duplication process will be deleted from the token stream, since a device connected to the output terminals (OUTDATA, OUTEXTN and OUTVALID) will not recognize these token words as valid data.

As before... and is duplicated.

Referring now more particularly to Figure 10, there is shown a reconfigurable process stage in accordance with one aspect of the present invention.

Input latches 34 receive an... the input latches 34 is passed as a first input over line 35 to a processing unit 36. A first output from the token decode subsystem 33 is passed over line 37 as a second input to the processing unit 36. A second output from the token decode 33 is passed over line 40 to an action identification

unit 39. The action identification unit 39 also receives input from registers 43 and 44 over line 46. The registers 43... is determined by the history of tokens previously received. The output from the action identification unit 39 is passed over line 38 as a third input to the processing unit 36. The output from the processing unit 36 is passed to output latches 41. The output from the output latches 41 is... decoder 56 is passed over line 63 as an input to an Index to Data Unit (ITOD) 64. The Huffman decoder 56 and the ITOD 64 work together as a single logical unit. The output from the ITOD 64 is passed over line 65 to an arithmetic logic unit (ALU) 66. A first output from the ALU 66 is passed over line 67 to... blocks 133.

Referring to Figure 14b, in the JPEG and H.261 standards, the Common Intermediate Format (CIF) is used, wherein a picture 141 is encoded as 6 rows each containing... in a zigzag direction indicated by the arrow 144. The GOBs 142 are, in turn, processed row-by-row, left-to-right in each row.

Referring now to Figure 14c, it... in accordance with the practice of the present invention. A first picture 161 to be processed contains a first PICTURE(underscore)START token 162, first-picture information of indeterminate length 163, and a first PICTURE(underscore)END token 164. A second picture 165 to be processed contains a second PICTURE(underscore)START token 166, second picture information of indeterminate length 167... tokens 162 and 166 indicate the start of the pictures 161 and 165 to the processor. Likewise, the PICTURE(underscore)END tokens 164 and 168 signify the end of the pictures 161 and 165 to the processor. This allows the processor to process picture information 163 and 167 of variable lengths.

Referring to Figure 17, a split 171...Video Formatter (not shown in Figure 17).

Referring now to Figure 18, the prediction filtering process is illustrated. A forward picture 201 is passed over line 202 as a first input... the right of the value decode shift register 230, as indicated by area 231. This process eliminates overlapping start code images, as discussed below. A first output from the value decode... Code Detector. The Start Code Detector then receives a first data value image 244. Before processing the first data value image 244, the Start Code Detector may detect a second start... image 244 at a length 246. If this occurs, the Start Code Detector does not process the first data value image 244, and instead receives and processes a second data value image 247.

Referring now ...line 1 of Table 600, whenever a "sequence start" image is received during H.261 processing or a "picture start" image is received during MPEG processing, the entire group of four control tokens is generated, each followed by its corresponding data...Picture Decoding

3. Motion Picture Decompression
4. RAM Memory Map
5. Bitstream Characteristics
6. Reconfigurable Processing Stage
7. Multi-Standard Coding
8. Multi-Standard Processing Circuit-2nd Mode of Operation
9. Start Code Detector
10. Tokens
11. DRAM Interface

12... described herein in greater detail) and reformatting this output for use, including display in a computer or other display systems, including a video display system. Implementation of this formatting varies significantly... the Spatial Decoder circuits.

The Spatial Decoder of the present invention performs all the required processing within a single picture. This reduces the redundancy within one picture.

The Temporal Decoder reduces... modeller 75, the inverse zig-zag 81 and the inverse DCT 83. The standard independent units within the Huffman decoder and parser include the ALU 66 and the token formatter 71.

Referring now to Figure 12, the standard-independent units include the DRAM interface 100, the fork 91, the FIFO register 96, the summer 98 and the output selector 106. The standard dependent units are the address generator 94,

which is different in H.261 and in MPEG, and...much of the operation is very similar between the three different compression standards.

The next unit is the state machine 68 (Figure 11) located within the Huffman decoder and parser. Here... ..The same holds true for JPEG, which is a third,completely independent program.

The next unit is the Huffman decoder 56 which functions with the index to data unit 64. Those two units cooperate together to perform the Huffman decoding. Here, the algorithm that is used for Huffman... ..to the Huffman decoder at different times consistent with the standard in operation.

The last unit on the chip that is dependent on the compression standard is the inverse quantizer 79...an H.261 group of blocks and an MPEG slice. When H.261 data is processed after the Start Code Detector, each group of blocks is preceded by a slice(underscore... ..these standards have totally different sets of tables.

As previously indicated, most of the system units are compression standard independent. If a unit is standard independent, and such units need not remember what CODING(underscore)STANDARD is being processed. All of the units that are standard dependent remember the compression standard as the CODING(underscore)STANDARD token flows... ..CODING(underscore)STANDARD tokens at the Start Code Detector that is positioned as the first unit in the pipeline, this change of compression standard is readily handled. The token says a... ..found in the standard, i.e. from the bitstream into a prediction mode token. This processing is performed by the Huffman decoder and parser state machine, where it is easy to... ..to that token. By having these tokens and using them appropriately, the design of other units in the machine is simplified. Although there may be some complications in the program, benefits... ..a first encoded signal (the MPEG or H.261 encoded video signal) in a pipeline processing system. The Temporal Decoder is not needed for JPEG decoding.

In this regard, the invention... ..the use of a single pipeline decoder and decompression system. The decoding and decompression pipeline processor is organized on a unique and special configuration which allows the handling of the multi... ..video signals through the use of techniques all compatible with the single pipeline decoder and processing system. The Spatial Decoder is combined with the Temporal Decoder, and the Video Formatter is...with only still pictures. The compression standard independent Spatial Decoder performs all of the data processing within the boundaries of a single picture. Such a decoder handles the spatial decompression of... ..to the multi-standard, configurable Video Formatter, which then provides an output to the display terminal. In a first sequence of similar pictures, each decompressed picture at the output of the... ..of control tokens and DATA tokens, in combination with a plurality of sequentially-positioned reconfigurable processing stages selected and organized to act as a standard-independent, reconfigurable-pipeline-processor.

With regard to JPEG decoding, a single Spatial Decoder with no off chip DRAM can... ..video. Accordingly, signals carried by DATA tokens pass directly through the Temporal Decoder without further processing when the Temporal Decoder is configured for a JPEG operation.

Another aspect of the present... ..for subsequent use in temporal decoding of subsequent pictures.

Generally, the Temporal Decoder performs the processing between pictures either earlier and/or later in time with reference to the picture currently... ..is distributed among several areas of DRAM in the sense that the decompressed output information, processed by the Spatial Decoder, is stored in other DRAM registers by other random access memories...first decoder circuit (the Spatial Decoder) directly to the Video Formatter for handling without signal processing delay.

The Temporal Decoder also reorders the blocks of picture data for display by a... ..from a selection of pictures which have arrived earlier or later than the picture under processing. When a picture is described in this context, it may mean any one of the... ..2. The result, i.e., the final decoded picture resulting from the addition of a process step performed by the decoder;

3. Previously decoded pictures read from the DRAM; and

4... ...START token and a subsequent PICTURE(underscore)END token.

After the picture data information is processed by the Temporal Decoder, it is either displayed or written back into a picture memory location. This information is then kept for further reference to be used in processing another different coded data picture.

Re-ordering of the MPEG encoded pictures for visual display...used to encode a referenced picture of a picture might be identified as being one unit long, another picture might be a number of units long, while still a third picture could be a fraction of that unit.

None of the existing standards (MPEG 1.2, JPEG, H.261) define a way of... ...picture rate, whereas the Video Formatter can handle a variable input picture rate.

6. RECONFIGURABLE PROCESSING STAGE

Referring again to Figure 10, the reconfigurable processing stage (RPS) comprises a token decode circuit 33 which is employed to receive the tokens... ...input latches 34. The output of the token decode circuit 33 is applied to a processing unit 36 over the two-wire interface 37 and an action identification circuit 39. The processing unit 36 is suitable for processing data under the control of the action identification circuit 39. After the processing is completed, the processing unit 36 connects such completed signals to the output, two-wire interface bus 40 through output... ...token decode circuit 33 are applied simultaneously to the action identification circuit 39 and the processing unit 36. The action identification function as well as the RPS is described in further detail... ...not standard independent circuits. The data flows through the token decode circuit 33, through the

processing unit 36 and onto the two-wire interface circuit 42 through the output latches 41. If... ...wire interface 42 through the output circuit 41. The present invention operates as a pipeline processor having a two-wire interface for controlling the movement of control tokens through the pipeline... ...time, the token decode circuit 33 provides a proper flag or index signal to the processing unit 36 to alert it to the presence of the token being handled by the action identification circuit 39.

Control tokens may also be processed.

A more detailed description of the various types of tokens usable in the present invention...standard now passing through the state machine shown with reference to Figure 10.

Similarly, the processing unit 36 which is under the control of the action identification circuit 39 is now ready to process the information contained in the data fields of the DATA token when it is appropriate... ...action identification circuit 39 and is immediately followed by a DATA token which is then processed by the processing unit 36. The control token exits the output latches circuit 41 over the output two-wire interface 42 immediately preceding the DATA token which has been processed within the processing unit 36.

In the present invention, the action identification circuit, 39, is a state machine holding... ...show that the action can also be affected by the token that is currently being processed by the token decode circuit 33.

In general, there is shown token decoding and data processing in accordance with the present invention. The data processing is performed as configured by the action identification circuit 39. The action is affected by... ...information stored from previously decoded tokens in registers 43 and 44, the current token under processing, and the state and history information that the action identification unit 39 has itself acquired. A distinction is thereby shown between Control tokens and DATA tokens.

In any RPS, some tokens are viewed by that RPS unit as being Control tokens in that they affect the operation of the RPS presumably at... ..are viewed by the RPS as DATA tokens. Such DATA tokens contain information which is processed by the RPS in a way that is determined by the design of the particular... ..view of the same token. Some of the tokens might be viewed by one RPS unit as DATA Tokens while another RPS unit might decide that it is actually a Control Token. For example, the quantization table information... ..into a token called a quantization table token (QUANT(underscore)TABLE) which goes down the processing pipeline. As far as that machine is concerned, all of that was data; it was... ..sort of data into another sort of data, which is clearly a function of the processing performed by that portion of the machine. However, when that information gets to the inverse... ..present. This information is viewed as control information, and then that control information affects the processing that is done on subsequent DATA tokens because it affects the number that you multiply... ..important feature of the invention is that each of the stages of circuitry has the processing capability within it to be able to perform the necessary operations for each of the... ..operations are to be performed at a given time, come as tokens. There is one processing element that differs between the different stages to provide this capability. In the state machine...standard is and it looks up the parameters that it needs to apply to the processing elements in order to perform a proper operation. For example, the inverse quantizer will look... ..is set to 1 for a particular compression standard, and will apply that to its processing circuitry.

In a similar sense the Huffman decoder 56 has a number of tables within... ..MPEG video standard or the JPEG video standard. These three compression coding standards specify similar processes to be done on the arriving data, but the structure of the datastreams is different... ..token stream embodying the current coding standard. The control tokens are passed through the pipeline processor, and are used, i.e., decoded, in the state machines to which they are relevant... ..this regard, the DATA Tokens are treated in the same fashion, insofar as they are processed only in the state machines that are configurable by the control tokens into processing such DATA Tokens. In the remaining state machines, they pass through unchanged.

More specifically, a... ..signals. The remaining portions of the token are used to indicate and identify the internal processing control function which is standard for all of the datastreams passing through the pipeline processor. In one form of the invention, the token extension is used to carry the current...accompanying data. As previously discussed, this information is utilized in the system to reconfigure the processing stage used to perform the function required by the various standards created for that purpose... ..picture number as indicated by the value.

The system also includes a multi-stage parallel processing pipeline operating under the principles of the two-wire interface previously described. Each of the... ..the token presently entering the state machine into the action identification circuit 39 or the processing unit 36, as appropriate. The processing unit has been previously reconfigured by the next previous control token into the form needed for handling the current coding standard, which is now entering the processing stage and carried by the next DATA token. Further, in accordance with this aspect of the invention, the succeeding state machines in the processing pipeline can be functioning under one coding standard, i.e., H.261, while a previous... ..tokens required to decode a number of coding standards with a fixed number of reconfigurable processing stages. More specifically, the PICTURE(underscore)END control token is employed because it is important... ..standard machine, it is necessary to create additional control tokens within the multi-standard pipeline processing machine which will then indicate which one of the standard decoding techniques to use. Such... ..and to push the current picture through the decoder to the display.

8. MULTI-STANDARD PROCESSING CIRCUIT - SECOND MODE OF OPERATION

A compression standard-dependent circuit, in the form of the...of the Start Code Detector will subsequently be discussed in further detail, as will the process of starting up of the decoder.

The aforementioned description has been concerned primarily with the ...the data which immediately follows according to the standard. However, in the multi-standard pipeline processing system of the present invention, where compatibility is required for multiple standards, the system has... ..signals, including flag signals, are

generated by each state machine to handle some of the processing within that state machine. Values carried in the standards can be used to access machine... its contents must be removed from the two wire interface to ensure that no further processing takes place using these 3 bytes. The decode register is emptied, and the value decode... 10. TOKENS

In the practice of the present invention, a token is a universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions and is adapted for use with a reconfigurable processing stage (RPS) which is a stage, which in response to a recognized token, reconfigures itself to perform various operations.

Tokens may be either position dependent or position independent upon the processing stages for performance of various functions. Tokens may also be metamorphic in that they can be altered by a processing stage and then passed down the pipeline for performance of further functions. Tokens may interact... other functions, and the specific interaction with a stage may be conditioned by the previous processing history of a stage.

A PICTURE(underscore)END token is a way of signalling the... through a fixed size, fixed width buffer.

The present invention is directed to a pipeline processing system which has a variable configuration which uses tokens and a two-wire system. The... do not use control tokens.

The control tokens are generated by circuitry within the decoder processor and emulate the operation of a number of different type standard-dependent signals passing into the serial pipeline processor for handling. The technique used is to study all the parameters of the multi-standards that are selected for processing by the serial processor and noting 1) their similarities, 2) their dissimilarities, 3) their needs and requirements and 4) selecting the correct token function to effectively process all of the standard signals sent into the serial processor. The functions of the tokens are to emulate the standards. A control token function is... the standard dependent signals and as an element to transmit control information through the pipeline processor.

In prior art system, a dedicated machine is designed according to well-known techniques to... tokens provide and make a sensible format for communicating information through the decompression circuit pipeline processor. In the design selected hereinafter and used in the preferred embodiment, each word of a... However, this is not a limitation on the invention, but on the magnitude of the processing steps elected to be accomplished by use of these tokens. It is to be noted... bit address for use in accessing the random access memories used throughout this serial decompression processor. This provides an additional degree of variability that facilitates a broad range of versatility.

As previously described, the DATA token carries data from one processing stage to the next. Consequently, the characteristics of this token change as it passes through... longest number of data bits because it needs to provide the most information to the

processing unit so that it can start the decompression with as much information as possible. Words which... to receive an address, it waits for the address generator to supply a valid address, processes that address and then sets the accept line high for one clock period. Thus, it... be read. This signal passes between two asynchronous clock regimes and, therefore, passes through three synchronizing flip flops.

Provided RAM2 312 is empty, the next item of data to arrive on... interesting.

In general, prediction data will be offset from the position of the block being processed as specified in the motion vectors in x and y. Thus, the block of data... address, 9. Data is read from this address and the x value is incremented. The process is repeated until the x value reaches its stop value, at which point, the y... is read, the x value is again incremented until it reaches its stop value. The process is repeated until both x and y values have reached their stop values. Thus, the... invention, is that additional information must be provided to the prediction filters to indicate what processing is required on the data. This consists of the following:

a "last byte" signal indicating... bit 0) is incremented and the x address (3 LSBS) is reset to zero. This process is repeated until 64 bytes have been read. With a 16 or 32 bit wide...register while its access register is set to zero, the results are undefined.

14. MICRO-PROCESSOR INTERFACE

A standard byte wide micro-processor interface (MPI) is used on all circuits with in the Spatial Decoder and Temporal Decoder... the parameter column. The actual specifications are shown in the respective columns min, max and units.

The DC operating conditions can be seen with reference to Table A.6.3. Here... the signal is present the maximum amount of time that this signal is available. The Units column gives the units of measurement used to describe the signals.

16. MPI WRITE TIMING

The general description of...a PICTURE(underscore)END token is decoded and forces the data in the coded data buffers to be applied to the Huffman decoder and video demultiplexor, the final picture can be... Consequently, the machine will not go into error recovery mode and will successfully continue to process the coded data.

A still further advantage of the use of a PICTURE(underscore)END token is that the serial pipeline processor will continue the processing of uninterrupted data. Through the use of a PICTURE(underscore)END token, the serial pipeline processor is configured to handle less than the expected amount of data and, therefore, continues processing. Typically, a prior art machine would stop itself because of an error condition. As previously... of the Huffman decode and Video Demultiplexor know the number of blocks that it will process during each picture recovery cycle. When the correct number of blocks do not arrive from... Each of the state machines recognizes a FLUSH control token as information not to be processed. Accordingly, the FLUSH token is used to fill up all of the remaining empty parts...less information than normally expected to decode the last picture. The Huffman decode circuit finishes processing the information contained in the last picture, and outputs this information through the DRAM interface... token, in accordance with the present invention, is used to pass through the entire pipeline processor and to ensure that the buffers are emptied and that other circuits are reconfigured to... underscore)END token, a padding word and a FLUSH token indicating to the serial pipeline processor that the picture processing for the current picture form is completed. Thereafter, the various state machines need reconfiguring to... FLUSH token resets each stage as it passes through, but-allows subsequent stages to continue processing. This prevents a loss of data. In other words, the FLUSH token is a variable... AFTER PICTURE

The STOP(underscore)AFTER(underscore)PICTURE function is employed to shut down the processing of the serial pipeline decompressing circuit at a logical point in its operation. At this... a picture, the STOP(underscore)AFTER(underscore)PICTURE operation signals the end of all current processing.

22. MULTI(underscore)STANDARD - SEARCH MODE

Another feature of the present invention is the use... underscore)MODE control token which is used to reconfigure the input to the serial pipeline processor to look at the incoming bit stream. When the search mode is set, the Start... combination of control tokens, and DATA tokens along with the reconfiguration circuits, to provide similar processing.

The use of search mode in the present invention is convenient in many situations including... video disc. In general, a search mode is convenient when the user interrupts the normal processing of the serial pipeline at a point where the machine does not expect such an...be the case.

In brief, the Huffman Decoder 321 works in conjunction with the other units shown in Figure 27. These other units are the Parser State Machine 322, the inshifter 323, the Index to Data unit 324, the ALU 325, and the Token Formatter 326. As described previously, connection between these blocks is governed by a two wire interface. A

more detailed description of how these units function is subsequently described herein in greater detail, the focus here is on particular aspects... control certain functions of the Index to Data 324 and ALU 325. Control of these units by the Huffman Decoder is necessary for proper decoding of block-level information. Having the... further detail in the "More Detailed Description of the Invention" section.

The Index to Data unit 324 performs the second part of the multi-part algorithm. This unit contains a look up table that provides the actual Huffman decoded data. Entries in the... by detecting these in the Huffman Decoder 321, rather than in the Index to Data unit 324.

This index number is then passed to the Index to Data unit 324. In essence, the Index to Data unit is a look-up table. In accordance with one aspect of the algorithm, the look... format that JPEG specifies for transferring an alternate JPEG table.

From the Index to Data unit 324, the decoded index number or other data is passed, together with the accompanying control... the entering data to ensure that the DATA tokens are of the correct size for processing. In fact, the token stream can be corrected in some situations if the error is... an order that is useful for the decompression circuits, but not for the particular display unit being used. When a block of data enters the Buffer Manager, the Buffer Manager supplies... the output of the Spatial Decoder or Temporal Decoder and re-format it for a computer or display system. The details of this formatting will vary between applications. In a simple... Token. The DATA Token can have as many bits as are necessary for carrying out processing at a particular place in the system. All other Tokens ignore the extra bits.

A.3.2 The DATA Token

The DATA Token carries data from one processing stage to the next. Consequently, the characteristics of this Token change as it passes through... will be sufficient to collect DATA Tokens and to detect a few Tokens that provide synchronization information (such as PICTURE(underscore)START). In this regard, see subsequent sections A.16, "Connecting... from the data stream. This provides an alternative to doing the configuration via the micro processor interface.

A.3.4 Description of Tokens

This section documents the Tokens which are implemented... 3.5.1. Note: JPEG requires a 2:1:1 structure for its macroblocks when processing 4:2:2 data. See Table A.3.5.

A.3.6 Special Token formats... either is low then the interface is taken to high impedance.

Note: on-chip data processing is not terminated when the DRAM interface is at high impedance. Therefore, errors will occur... decoded video's picture rate. Accordingly, this clock can be used to provide audio/video synchronization.

A.7.1 Spatial Decoder clock signals

The Spatial Decoder has two different (and potentially... in accordance with the present invention, must know what video standard is being input for processing. Thereafter, the system can accept either pre-existing Tokens or raw byte data which is... time a value is written into coded(underscore)data (7:0). Software is responsible for settling coded(underscore)extn to 0 before the last word of any Token is written to... 0). The start of this new DATA Token then passes into the Spatial Decoder for processing.

Each time a new 8 bit value is written to coded(underscore)data (7:0... Detector analyses data in the DATA Tokens bit serially. The Detector's normal rate of processing is one bit per clock cycle (of coded(underscore)clock). Accordingly, it will typically decode a byte of coded data every 8 cycles of coded(underscore)clock. However, extra processing cycles are occasionally required, e.g., when a non-DATA Token is supplied or when main decoder(underscore)clock. Data transfer is synchronized to decoder(underscore)clock on-chip.

SECTION A.11 Start code detector

A.11.1... ..Code Detector. So, accessing these registers will be unreliable if the Start Code Detector is processing data. The user is responsible for ensuring that the Start Code Detector is halted before... ..Detector. In this case, the Tokens are passed through the Start Code Detector with no processing to other stages of the Spatial Decoder. These Tokens can only be inserted just before...result will be unpredictable if this is done when the Start Code Detector is actively processing data.

Discard all mode can be safely initiated after any of the Start Code Detector...start code non-alignment interrupt is suppressed.

In contrast, however, JPEG was designed for a computer environment where byte alignment is guaranteed. Therefore, marker codes should only be detected when byte... ..the other hand, was designed to meet the needs of both communications (bit serial) and

computer (byte oriented) systems. Start codes in MPEG data should normally be byte aligned. However, the...result will be unpredictable if this is done when the Start Code Detector is actively processing data. So, before initiating a start code search, the Start Code Detector should be stopped so no data is being processed. The Start Code Detector is always in this condition if any of the Start Code...the spatial video decoding circuits (inverse modeler, quantizer and DCT). This second logical buffer allows processing time to include a spread so as to accommodate processing pictures having varying amounts of data.

Both buffers are physically held in a single off the unit for all the above mentioned registers is a 512 bit block of data. Accordingly, the... ..until there is space in the buffer. If a buffer continues to be full, more processing stages "up steam" of the buffer will halt until the Spatial Decoder is unable to... ..converting coded data into Tokens started by the Start Code Detector. There are four main processing blocks in the Video Demux: Parser State Machine, Huffman decoder (including an ITOD), Macroblock counter... ..or state machine follows the syntax of the coded video data and instructs the other units. The Huffman decoder converts variable length coded (VLC) data into integers. The Macroblock counter keeps...

Claims:

1. A system having a plurality of processing stages, comprising a universal adaptation unit in the form of an interactive interfacing token for control and/or data functions among said processing stages,

wherein said token is a CODING(underscore)STANDARD token for conditioning said system for processing in a selected one of a plurality of picture compression/ decompression standards; one of said processing stages being a Huffman decoder and parser; one of said control tokens being a CODING... ..Data stage, and wherein said parser stage sends an instruction to said Index to Data Unit to select tables needed for a particular identified coding standard, said parser stage indicating whether...

6/K/5 (Item 5 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Country	Number	Kind	Date		
Type		Pub. Date		Kind	Text
Available Text		Language		Update	Word Count
Total Word Count (Document A)					
Total Word Count (Document B)					
Total Word Count (All Documents)					

Specification: INTRODUCTION

The present invention is directed to improvements in methods and apparatus for decompression which operates to decompress and/or decode a plurality of differently encoded input... ..of the well known standards known as JPEG, MPEG and H.261.

A serial pipeline processing system of the present invention comprises a single two-wire bus used for carrying... ..to a plurality of adaptive decompression circuits and the like positioned as a reconfigurable pipeline processor.

PRIOR ART

One prior art system is described in United States Patent No. 5,216,724. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and a host processor. The device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus.

United States Patent No. 4,785... ..a known quad tree data structure.

United States Patent No. 5,122,875 discloses an apparatus for encoding/decoding an HDTV signal. The apparatus includes a compression circuit responsive to high definition video source signals for providing hierarchically layered... ..compressed video data of relatively greater and lesser importance to image reproduction respectively. A transport processor, responsive to the high and low priority codeword sequences, forms high and low priority transport... ..United States Patent No. 5,168,356 discloses a video signal encoding system that includes apparatus for segmenting encoded video data into transport blocks for signal ...in respective transport blocks.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening. This is accomplished by an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ..frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ..oscillators and the receiver can continuously receive each channel, then the receiver need not be synchronized with the transmitter. An FFT algorithm implements a fast discrete approximation to the continuous case in which the receiver synchronizes to the first frame and then acquires subsequent frames every frame period. The frame period... ..increasing the amount of data transmitted.

United States Patent No. 5,212,742 discloses an apparatus and method for processing video data for compression/decompression in real-time. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and host processor. Lastly, the device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus. The method handles assigning portions of the image for each of the processors to operate upon.

United States Patent No. 5,231,484 discloses a system and method... ..MPEG standards. Included are three cooperating components or subsystems that operate to variously adaptively pre-process the incoming digital motion video sequences, allocate bits to the pictures in a sequence, and... ..States Patent No. 5,267,334 discloses a method of removing frame redundancy in a computer system for a sequence of moving images. The method comprises detecting a first scene change... ..facing" keyframe or intraframe, and it is normally present in CCITT compressed video data. The process then comprises generating at least one intermediate compressed frame, the at least one intermediate compressed frame containing difference information from the first image for at least one image following...change, known as a "backward-facing" keyframe. The first keyframe and the at least one intermediate compressed frame are linked for forward play, and the second keyframe and the intermediate compressed frames are linked in reverse for reverse play. The intraframe may also be used... ..of complete scene information.

United States Patent No. 5,276,513 discloses a first circuit apparatus, comprising a given number of prior-art image-pyramid stages, together with a second circuit apparatus, comprising the same given number of novel motion-vector stages, perform cost-effective hierarchical motion analysis (HMA) in real-time, with minimum system processing delay and/or employing minimum system processing delay and/or employing minimum hardware structure. Specifically, the first and second circuit apparatus, in response to relatively high-resolution image data from an ongoing input series of successive... ..a relatively high frame rate (e.g., 30 frames per second), derives, after a certain processing-system delay, an ongoing output series of successive given pixel-density vector-data frames that... ..of successive image frames.

United States Patent No. 5,283,646 discloses a method and apparatus for enabling a real-time video encoding system to accurately deliver the desired number of... ..desired bit allocations.

The article, Chong, Yong M., A Data-Flow Architecture for Digital Image Processing, Wescon Technical Papers: No. 2 Oct./Nov. 1984, discloses a real-time signal processing system specifically designed for image processing. More particularly, a token based data-flow architecture is disclosed wherein the tokens are of... ..width having a fixed width address field. The system contains a plurality of identical flow processors connected in a ring fashion. The tokens contain a data field, a control field and a tag. The tag field of the token is further broken down into a processor address field and an identifier field. The processor address field is used to direct the tokens to the correct data-flow processor, and the identifier field is used to label the data such that the data-flow processor knows what to do with the data. In this way, the identifier field acts as an instruction for the data-flow processor. The system directs each token to a specific data-flow processor using a module number (MN). If the MN matches the MN of the particular stage... ..to locate the decoder in the preceding stage in order to pre-decode complex decoding processing and to alleviate critical path problems in the logic circuit. The elastic nature of the...of block signal in most cases.

United States Patent No. 4,903,018 discloses a process and data processing system for compressing and expanding structurally associated multiple data sequences. The process is particular to data sets in which an analysis is made of the structure in... ..data series on the basis of the order number of these data elements. The data processing system for performing the processes includes a storage matrix (26) and an index storage (28) having line addresses of the... ..the final actual video.

United States Patent No. 5,060,242 discloses an image signal processing system DPCM encodes the signal, then Huffman and run length encodes the signal to produce... ..tightly packed without gaps for efficient transmission without loss of any data. The tightly packed apparatus has a barrel shifter with its shift modulus controlled by an accumulator receiving code word... ..OR gate is connected to the shifter, while a register is connected to the gate. Apparatus for processing a tightly packed and decorrelated digital signal has a barrel shifter and accumulator for unpacking... ..an inverse DCPM decoder.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening is accomplished by use of an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed

by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ..frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ..kernel matrix.

United States Patent No. 5,231,486 discloses a high definition video system processes a bitstream including high and low priority variable length coded Data words. The coded Data... ..packed High Priority Data and packed Low Priority Data by means of respective data packing units. The coded Data is continuously applied to both packing units. High Priority and Low Priority Length words indicating the bit lengths of high priority and... ..States Patent No. 5,287,178 discloses a video signal encoding system includes a signal processor for segmenting encoded video data into transport blocks having a header section and a packed data section. The system also includes reset control apparatus for releasing resets of system components, after a global system reset, in a prescribed non-simultaneous phased sequence to enable signal processing to commence in the prescribed sequence. The phased reset release sequence begins when valid data...United States Patent No. 5,142,380 to Sakagami et al. discloses an image compression apparatus suitable for use with still images such as those formed by electronic still cameras using... ..and Q.

United States Patent No. 5,193,002 to Guichard et al. disclosed an apparatus for coding/decoding image signals in real time in conjunction with the CCITT standard H.261. A digital signal processor carries out direct quantization and reverse quantization.

United States Patent No. 5,241,383 to Chen et al. describes an apparatus with a pseudo-constant bit rate video coding achieved by an adjustable quantization parameter. The... ..relates to an improved pipeline system having an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along the pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the processing stages in the pipeline and interacting with selected stages in the pipeline for control data and/or combined control-data functions among the processing stages, so that the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing. In accordance with the invention, the processing stages may be configurable in response to recognition of at least one token. One of the processing stages may be a Start Code Detector which receives the input and generates and/or... ..and resetting the system, and a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards. The present invention... ..data and having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video pictures of varying data size can be controlled. Also in accordance with the invention, a processing stage receives the input data stream, the stage including means for recognizing specified bit stream patterns, whereby the processing stage facilitates random access and error recovery. The invention may also include a means for... ..invention also includes an inverse modeller stage, an inverse discrete cosine transform stage, and a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for processing data.

In addition, the present invention relates to an improved pipeline system having a Huffman...pipeline stage that incorporates a two-wire transfer control and also shows two consecutive pipeline processing stages with the two-wire transfer control;

Figures. 5a and 5b taken together depict one... ..shown in Figures. 8a and 8b.

Figure 10 is a block diagram of a reconfigurable processing stage;

Figure 11 is a block diagram of a spatial decoder;

Figure 12 is a... ..decoder including the prediction filters;

Figure 18 is a pictorial representation of the prediction filtering process;

Figure 19 shows a generalized representation of the macroblock structure;

Figure 20 shows a generalized... ..buffer;

Figure 25 is a pictorial diagram illustrating prediction data offset from the block being processed;

Figure 26 is a pictorial diagram illustrating prediction data offset by (1,1);

Figure 27...in general terms, the present invention provides an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along a pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the stages in the pipeline and interacting with selected stages in the pipeline for control, data and/or combined control-data functions among the processing stages, whereby the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing.

Each of the processing stages in the pipeline may include both primary and secondary storage, and the stages in processing stages for performance of functions or position independent of the processing stages for performance of functions.

In a pipeline machine, in accordance with the invention, the... ..altered by interfacing with the stages, and the tokens may interact with all of the processing stages in the pipeline or only with some but less than all of said processing stages. The tokens in the pipeline may interact with adjacent processing stages or with non-adjacent processing stages, and the tokens may reconfigure the processing stages. Such tokens may be position dependent for some functions and position independent for other... ..be Huffman coded.

In the improved pipeline machine, the tokens may be generated by a processing stage. Such pipeline tokens may include data for transfer to the processing stages or the tokens may be devoid of data. Some of the tokens may be identified as DATA tokens and provide data to the processing stages in the pipeline, while other tokens are identified as control tokens and only condition the processing stages in the pipeline, such conditioning including reconfiguring of the processing stages. Still other tokens may provide both data and conditioning to the processing stages in the pipeline. Some of said tokens may identify coding standards to the processing stages in the pipeline, whereas other tokens may operate independent of any coding standard among the processing stages. The tokens may be capable of successive alteration by the processing stages in the pipeline.

In accordance with the invention, the interactive flexibility of the tokens in cooperation with the processing stages facilitates greater functional diversity of the processing stages for resident structure in the pipeline, and the flexibility of the tokens facilitates system... ..or alteration. The tokens may be capable of facilitating a plurality of functions within any processing stage in the pipeline. Such pipeline tokens may be either hardware based or software based... ..system bandwidth in the pipeline. The tokens may provide data and control simultaneously to the processing stages in the pipeline.

The invention may include a pipeline processing machine for handling plurality of separately encoded bit streams arranged as a single serial bit... ..and for passing unrecognized control tokens along the pipeline, and a reconfigurable decode and parser processing means responsive to a recognized control token for reconfiguring a particular stage to handle an... ..be a pipeline system and the Start Code Detector may be positioned as the first processing stage in the pipeline.

The present invention also provides, in a system having a plurality of processing stages, a universal adaptation unit in the form of an interactive interfacing token for control and/or data functions among the processing stages, the token being a PICTURE(underscore)START code token for indicating that the start... ..The token may also be a

CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards.

The CODING(underscore... standard as JPEG, and/or any other appropriate picture standard. At least some of the processing stages reconfigure in response to the CODING(underscore)STANDARD token.

One of the processing stages in the system may be a Huffman decoder and parser and, upon receipt of... Data stage, and the parser stage may send an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... video data, having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video... controlled.

The system may include a spatial decoder having a two-wire interface interconnecting processing stages, the interface enabling serial processing for data and parallel processing for control.

As previously indicated, the system may further include a ROM having separate stored... of a plurality of picture standards, the programs being selectable by a token to facilitate processing for a plurality of different picture standards.

The spatial decoder system also includes a token... decoding stage and a parser stage for sending an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... The present invention also provides a pipeline system having an input data stream, and a processing stage for receiving the input data stream, the stage including means for recognizing specified bit... whereby said stage facilitates random access and error recovery. In accordance with the invention, the processing stage may be a start code detector and the bit stream patterns may include start... token and padding insures uniformity of word size. In accordance with the invention, a reconfigurable processing stage may be provided as a spatial decoder and the padding means adds to picture... that if the DATA token has less than the predetermined length, the padder circuit adds units of data to the DATA token until the predetermined length is achieved. A bypass circuit... tokens into a buffer, having a second predetermined width.

The invention also provides an apparatus for providing a time delay to a group of compressed pictures, the pictures corresponding to... and capable of delaying the words of data, is in communication with a control circuit intermediate the counter circuit and the inverse modeller circuit, the control circuit also communicating with the... inverse modeller stage and an inverse discrete cosine transform stage, the improvement characterized by a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for processing data.

In accordance with the invention, the token may be a QUANT(underscore)TABLE token for causing the

processing stage to generate a quantization table.

The present invention also provides a Huffman decoder for... of bits used to represent an item of data.

DECODER: An embodiment of a decoding process.

DECODING (PROCESS): The process defined in this specification that reads an input coded bitstream and produces decoded pictures or... the same order in which they were presented at the input of the encoder.

ENCODING (PROCESS): A process, not specified in this specification, that reads a stream of input pictures or audio samples... to provide an estimate of the pel value or data element currently being decoded.

RECONFIGURABLE PROCESS STAGE (RPS): A stage, which in response to a recognized token, reconfigures itself to perform various operations.

SLICE: A series of macroblocks.

TOKEN: A universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions... ..indicates that the corresponding stage holds valid data, i.e., data that is to be processed in one of the pipeline stages. After processing (which may involve nothing more than a simple transfer without manipulation of the data) valid... ..present invention may be used with any number of pipeline stages. Furthermore, data may be processed in more than one stage and the processing time for different stages can differ.

In addition to clock and data signals (described below... ..other system. For example, the last pipeline stage may pass its data on to subsequent processing circuitry. The ACCEPT signal, which is illustrated as the lower of the two lines connecting... ..the minimum disturbance possible to other pipeline stages. Succeeding pipeline stages are allowed to continue processing and, therefore, this means that gaps open up in the stream of data following the... ..The data in the pipeline is encoded such that many different types of data are processed in the pipeline. This encoding accommodates data packets of variable size and the size of...the other hand, it may generate itself, all or part of the data to be processed in the pipeline. Indeed, as is explained below, a "stage" may contain arbitrary processing circuitry, including none at all (for simple passing of data) or entire systems (for example... ..values zero and 255 may not be used.

If such a picture were to be processed in a pipeline built in the practice of the present invention, then one of these... ..data must not be written over since it is data that must be saved for processing or use in a downstream device e.g., a pipeline stage, a device or a... ..connected to the pipeline upstream contains data D4 that is to be transferred into and processed in the pipeline. ...pipeline, in accordance with the preferred embodiments of the present invention, to "fill up" empty processing stages is highly advantageous since the processing stages in the pipeline thereby become decouple from one another. In other words, even though... ..data can be transferred into the pipeline and between stages even when one or more processing stages is blocked.

In the embodiment shown in Fig. 1, it is assumed that the...propagate all the way back to the beginning of the pipeline if there is some intermediate stage that is able to accept new data.

In the embodiment illustrated in Fig. 1...has been mentioned. It is to be further understood that each pipeline stage may also process the data it has received arbitrarily before passing it between its internal storage elements or... ..the portion of the pipeline that contains input and output storage elements and that arbitrarily processes data stored in its storage elements.

Furthermore, the "device" downstream from ...valid data, but also when a stage requires more than one clock phase to finish processing its data. This also can occur when it creates valid data in one or both... ..control the passage of data between adjacent storage elements. The VALID signal may also be processed in an analogous manner.

A great advantage of the two-wire interface (one wire for... ..In addition, two extra latches and a small number of gates are preferably added to process the ACCEPT and VALID signals that are associated with the data latches in each half... ..application so requires. The interface in accordance with this embodiment can also be used to process analog signals.

As discussed previously, while other conventional timing arrangements may be used, the interface... ..circuit B1, which may be provided to convert output data from input latch LDIN into intermediate data, which is then later loaded in an output data latch LDOUT, which comprises the... ..is connected either directly as an input to the validation output latch LVOUT, or via intermediate logic devices or circuits that may alter the signal.

Similarly, the output validation signal QVOUT... ..to the input of the validation input latch QVIN of the following stage, or via intermediate devices or logic circuits, which may alter the validation signal. This output QVIN ...word.

Preferred Data Structure - "tokens"

In the sample application shown in Fig. 4, each stage processes all input data, since there is no control circuitry that excludes any stage from allowing... are connected together in a relatively simple configuration. The simplest configuration is a pipeline of processing steps. For example, in the one shown in Fig. 1. The use of tokens, however... flows from left to right in the diagram. Data enters the machine and passes into processing Stage A. This may or may not modify the data and it then passes the... advantage of the tokens is their ability to achieve this kind of communication. Since any processing stage that does not recognize a token simply passes it on unaltered to the next... is transmitted along with the address and data fields in each token so that a processing stage can pass on a token (which can be of arbitrary length) without having to... be the first word of a new token.

Note that although the simple pipeline of processing stages is particularly useful, it will be appreciated that tokens may be applied to more complicated configurations of processing elements. An example of a more complicated processing element is described below.

It is not necessary, in accordance with the present invention, to... has extension bits. An example of this is a token that activates a stage that processes video quantization values stored in a quantization table (typically a memory device). For example, a... turn, is of great importance in video data pipeline systems since it ensures that all processing stages can be continuously running at full bandwidth.

In accordance to the present invention, in... some other chips in the set. This is advantageous both from the perspective of a customer and from that of a chip manufacturer. Even if modifications mean that all chips are... the end of a token (and hence the start of the next token) to be processed correctly (including simple non-manipulative transfer), even if the token is not recognized by the... block diagram of a pipeline stage whose function is as follows. If the stage is processing a predetermined token (known in this example as the DATA token), then it will duplicate... the address field of the DATA token. If, on the other hand, the stage is processing any other kind of token, it will delete every word. The overall effect is that... respective output signals:

In the duplication stage, the output from the data latch LDIN forms intermediate data referred to as MID(underscore)DATA. This intermediate data word is loaded into the data output latch LDOUT only when an intermediate acceptance signal (labeled "MID(underscore)ACCEPT" in Fig. 8a) is set HIGH.

The portion of... data. These include a "DATA(underscore)TOKEN" signal that indicates that the circuitry is currently processing a valid DATA Token, and a NOT(underscore)DUPLICATE signal which is used to control duplication of data. When the circuitry is processing a DATA Token, the NOT(underscore)DUPLICATE signal toggles between a HIGH and a LOW... the token to be duplicated once (but no more times). When the circuitry is not processing a valid DATA Token then the NOT(underscore)DUPLICATE signal is held in a HIGH state. Accordingly, this means that the token words that are being processed are not duplicated.

As Fig. 8a illustrates, the upper six bits of 8-bit intermediate data word and the output signal QI1 from the latch LI1 form inputs to a... explained further below.

Latch LO1 performs the function of latching the last value of the intermediate extension bit (labeled "MID(underscore)EXTN" and as signal S4), and it loads this value... and the DATA(underscore)TOKEN signal will become "0", indicating that the circuitry is not processing a DATA token.

If QI1 is "0" and SO is "0", thereby indicating a DATA... phase and the DATA(underscore)TOKEN signal will become "1", indicating that the circuitry is processing a DATA token.

The NOT(underscore)DUPLICATE signal (the output signal QO3) is similarly loaded... LVOUT at the same time that MID(underscore)DATA is loaded into LDOUT and the intermediate extension bit (signal S4) is loaded into LEOUT. Signal S5 is also combined with the... above. This has the effect that all tokens except the one that causes the duplication process will be deleted from the token stream, since a device connected to the output terminals (OUTDATA, OUTEXTN and OUTVALID) will not recognize these token words as valid data.

As before... ..and is duplicated.

Referring now more particularly to Figure 10, there is shown a reconfigurable process stage in accordance with one aspect of the present invention.

Input latches 34 receive an... ..the input latches 34 is passed as a first input over line 35 to a processing unit 36. A first output from the token decode subsystem 33 is passed over line 37 as a second input to the processing unit 36. A second output from the token decode 33 is passed over line 40 to an action identification

unit 39. The action identification unit 39 also receives input from registers 43 and 44 over line 46. The registers 43... ..is determined by the history of tokens previously received. The output from the action identification unit 39 is passed over line 38 as a third input to the processing unit 36. The output from the processing unit 36 is passed to output latches 41. The output from the output latches 41 is... ..decoder 56 is passed over line 63 as an input to an Index to Data Unit (ITOD) 64. The Huffman decoder 56 and the ITOD 64 work together as a single logical unit. The output from the ITOD 64 is passed over line 65 to an arithmetic logic unit (ALU) 66. A first output from the ALU 66 is passed over line 67 to...blocks 133.

Referring to Figure 14b, in the JPEG and H.261 standards, the Common Intermediate Format (CIF) is used, wherein a picture 141 is encoded as 6 rows each containing... ..in a zigzag direction indicated by the arrow 144. The GOBs 142 are, in turn, processed row-by-row, left-to-right in each row.

Referring now to Figure 14c, it... ..in accordance with the practice of the present invention. A first picture 161 to be processed contains a first PICTURE(underscore)START token 162, first-picture information of indeterminate length 163, and a first PICTURE(underscore)END token 164. A second picture 165 to be processed contains a second PICTURE(underscore)START token 166, second picture information of indeterminate length 167... ..tokens 162 and 166 indicate the start of the pictures 161 and 165 to the processor. Likewise, the PICTURE(underscore)END tokens 164 and 168 signify the end of the pictures 161 and 165 to the processor. This allows the processor to process picture information 163 and 167 of variable lengths.

Referring to Figure 17, a split 171...Video Formatter (not shown in Figure 17).

Referring now to Figure 18, the prediction filtering process is illustrated. A forward picture 201 is passed over line 202 as a first input... ..the right of the value decode shift register 230, as indicated by area 231. This process eliminates overlapping start code images, as discussed below. A first output from the value decode... ..Code Detector. The Start Code Detector then receives a first data value image 244. Before processing the first data value image 244, the Start Code Detector may detect a second start... ..image 244 at a length 246. If this occurs, the Start Code Detector does not process the first data value image 244, and instead receives and processes a second data value image 247.

Referring now ...line 1 of Table 600, whenever a "sequence start" image is received during H.261 processing or a "picture start" image is received during MPEG processing, the entire group of four control tokens is generated, each followed by its corresponding data...Picture Decoding

3. Motion Picture Decompression
4. RAM Memory Map
5. Bitstream Characteristics
6. Reconfigurable Processing Stage
7. Multi-Standard Coding
8. Multi-Standard Processing Circuit-2nd Mode of Operation

9. Start Code Detector

10. Tokens

11. DRAM Interface

12... ..described herein in greater detail) and reformatting this output for use, including display in a computer or other display systems, including a video display system. Implementation of this formatting varies significantly... ..the Spatial Decoder circuits.

The Spatial Decoder of the present invention performs all the required processing within a single picture. This reduces the redundancy within one picture.

The Temporal Decoder reduces... ..modeller 75, the inverse zig-zag 81 and the inverse DCT 83. The standard independent units within the Huffman decoder and parser include the ALU 66 and the token formatter 71.

Referring now to Figure 12, the standard-independent units include the DRAM interface 100, the fork 91, the FIFO register 96, the summer 98 and the output selector 106. The standard dependent units are the address generator 94, which is different in H.261 and in MPEG, and...much of the operation is very similar between the three different compression standards.

The next unit is the state machine 68 (Figure 11) located within the Huffman decoder and parser. Here... ..The same holds true for JPEG, which is a third, completely independent program.

The next unit is the Huffman decoder 56 which functions with the index to data unit 64. Those two units cooperate together to perform the Huffman decoding. Here, the algorithm that is used for Huffman... ..to the Huffman decoder at different times consistent with the standard in operation.

The last unit on the chip that is dependent on the compression standard is the inverse quantizer 79...an H.261 group of blocks and an MPEG slice. When H.261 data is processed after the Start Code Detector, each group of blocks is preceded by a slice(underscore... ..these standards have totally different sets of tables.

As previously indicated, most of the system units are compression standard independent. If a unit is standard independent, and such units need not remember what CODING(underscore)STANDARD is being processed. All of the units that are standard dependent remember the compression standard as the CODING(underscore)STANDARD token flows... ..CODING(underscore)STANDARD tokens at the Start Code Detector that is positioned as the first unit in the pipeline, this change of compression standard is readily handled. The token says a... ..found in the standard, i.e. from the bitstream into a prediction mode token. This processing is performed by the Huffman decoder and parser state machine, where it is easy to... ..to that token. By having these tokens and using them appropriately, the design of other units in the machine is simplified. Although there may be some complications in the program, benefits... ..a first encoded signal (the MPEG or H.261 encoded video signal) in a pipeline processing system. The Temporal Decoder is not needed for JPEG decoding.

In this regard, the invention... ..the use of a single pipeline decoder and decompression system. The decoding and decompression pipeline processor is organized on a unique and special configuration which allows the handling of the multi... ..video signals through the use of techniques all compatible with the single pipeline decoder and processing system. The Spatial Decoder is combined with the Temporal Decoder, and the Video Formatter is...with only still pictures. The compression standard independent Spatial Decoder performs all of the data processing within the boundaries of a single picture. Such a decoder handles the spatial decompression of... ..to the multi-standard, configurable Video Formatter, which then provides an output to the display terminal. In a first sequence of similar pictures, each decompressed picture at the output of the... ..of control tokens and DATA tokens, in combination with a plurality of sequentially-positioned reconfigurable processing stages selected and organized to act as a standard-independent, reconfigurable-pipeline-processor.

With regard to JPEG decoding, a single Spatial Decoder with no off chip DRAM can... video. Accordingly, signals carried by DATA tokens pass directly through the Temporal Decoder without further processing when the Temporal Decoder is configured for a JPEG operation.

Another aspect of the present... for subsequent use in temporal decoding of subsequent pictures.

Generally, the Temporal Decoder performs the processing between pictures either earlier and/or later in time with reference to the picture currently... is distributed among several areas of DRAM in the sense that the decompressed output information, processed by the Spatial Decoder, is stored in other DRAM registers by other random access memories... first decoder circuit (the Spatial Decoder) directly to the Video Formatter for handling without signal processing delay.

The Temporal Decoder also reorders the blocks of picture data for display by a... from a selection of pictures which have arrived earlier or later than the picture under processing. When a picture is described in this context, it may mean any one of the... 2. The result, i.e., the final decoded picture resulting from the addition of a process step performed by the decoder;

3. Previously decoded pictures read from the DRAM; and

4... START token and a subsequent PICTURE(underscore)END token.

After the picture data information is processed by the Temporal Decoder, it is either displayed or written back into a picture memory location. This information is then kept for further reference to be used in processing another different coded data picture.

Re-ordering of the MPEG encoded pictures for visual display... used to encode a referenced picture of a picture might be identified as being one unit long, another picture might be a number of units long, while still a third picture could be a fraction of that unit.

None of the existing standards (MPEG 1.2, JPEG, H.261) define a way of... picture rate, whereas the Video Formatter can handle a variable input picture rate.

6. RECONFIGURABLE PROCESSING STAGE

Referring again to Figure 10, the reconfigurable processing stage (RPS) comprises a token decode circuit 33 which is employed to receive the tokens... input latches 34. The output of the token decode circuit 33 is applied to a processing unit 36 over the two-wire interface 37 and an action identification circuit 39. The processing unit 36 is suitable for processing data under the control of the action identification circuit 39. After the processing is completed, the processing unit 36 connects such completed signals to the output, two-wire interface bus 40 through output... token decode circuit 33 are applied simultaneously to the action identification circuit 39 and the processing unit 36. The action identification function as well as the RPS is described in further detail... not standard independent circuits. The data flows through the token decode circuit 33, through the

processing unit 36 and onto the two-wire interface circuit 42 through the output latches 41. If... wire interface 42 through the output circuit 41. The present invention operates as a pipeline processor having a two-wire interface for controlling the movement of control tokens through the pipeline... time, the token decode circuit 33 provides a proper flag or index signal to the processing unit 36 to alert it to the presence of the token being handled by the action identification circuit 39.

Control tokens may also be processed.

A more detailed description of the various types of tokens usable in the present invention... standard now passing through the state machine shown with reference to Figure 10.

Similarly, the processing unit 36 which is under the control of the action identification circuit 39 is now ready to process the information contained in the data fields of the DATA token when it is appropriate... action identification circuit 39 and is immediately followed by a DATA token which is then processed by the processing unit 36. The control token exits the output latches circuit 41 over the output two-wire interface 42 immediately preceding the DATA token which has been processed within the processing unit 36.

In the present invention, the action identification circuit, 39, is a state machine holding... show that the action can also be affected by the token that is currently being processed by the token decode circuit 33.

In general, there is shown token decoding and data processing in accordance with the present invention. The data processing is performed as configured by the action identification circuit 39. The action is affected by... information stored from previously decoded tokens in registers 43 and 44, the current token under processing, and the state and history information that the action identification unit 39 has itself acquired. A distinction is thereby shown between Control tokens and DATA tokens.

In any RPS, some tokens are viewed by that RPS unit as being Control tokens in that they affect the operation of the RPS presumably at... are viewed by the RPS as DATA tokens. Such DATA tokens contain information which is processed by the RPS in a way that is determined by the design of the particular... view of the same token. Some of the tokens might be viewed by one RPS unit as DATA Tokens while another RPS unit might decide that it is actually a Control Token. For example, the quantization table information... into a token called a quantization table token (QUANT(underscore)TABLE) which goes down the processing pipeline. As far as that machine is concerned, all of that was data; it was... sort of data into another sort of data, which is clearly a function of the processing performed by that portion of the machine. However, when that information gets to the inverse... present. This information is viewed as control information, and then that control information affects the processing that is done on subsequent DATA tokens because it affects the number that you multiply... important feature of the invention is that each of the stages of circuitry has the processing capability within it to be able to perform the necessary operations for each of the... operations are to be performed at a given time, come as tokens. There is one processing element that differs between the different stages to provide this capability. In the state machine... standard is and it looks up the parameters that it needs to apply to the processing elements in order to perform a proper operation. For example, the inverse quantizer will look... is set to 1 for a particular compression standard, and will apply that to its processing circuitry.

In a similar sense the Huffman decoder 56 has a number of tables within... MPEG video standard or the JPEG video standard. These three compression coding standards specify similar processes to be done on the arriving data, but the structure of the datastreams is different... token stream embodying the current coding standard. The control tokens are passed through the pipeline processor, and are used, i.e., decoded, in the state machines to which they are relevant... this regard, the DATA Tokens are treated in the same fashion, insofar as they are processed only in the state machines that are configurable by the control tokens into processing such DATA Tokens. In the remaining state machines, they pass through unchanged.

More specifically, a... signals. The remaining portions of the token are used to indicate and identify the internal processing control function which is standard for all of the datastreams passing through the pipeline processor. In one form of the invention, the token extension is used to carry the current... accompanying data. As previously discussed, this information is utilized in the system to reconfigure the processing stage used to perform the function required by the various standards created for that purpose... picture number as indicated by the value.

The system also includes a multi-stage parallel processing pipeline operating under the principles of the two-wire interface previously described. Each of the... the token presently entering the state machine into the action identification circuit 39 or the processing unit 36, as appropriate. The processing unit has been previously reconfigured by the next previous control token into the form needed for handling the current coding standard, which is now entering the processing stage and carried by the next DATA token. Further, in accordance with this aspect of

the invention, the succeeding state machines in the processing pipeline can be functioning under one coding standard, i.e., H.261, while a previous... tokens required to decode a number of coding standards with a fixed number of reconfigurable processing stages. More specifically, the PICTURE(underscore)END control token is employed because it is important... standard machine, it is necessary to create additional control tokens within the multi-standard pipeline processing machine which will then indicate which one of the standard decoding techniques to use. Such... and to push the current picture through the decoder to the display.

8. MULTI-STANDARD PROCESSING CIRCUIT - SECOND MODE OF OPERATION

A compression standard-dependent circuit, in the form of the...of the Start Code Detector will subsequently be discussed in further detail, as will the process of starting up of the decoder.

The aforementioned description has been concerned primarily with the ...the data which immediately follows according to the standard. However, in the multi-standard pipeline processing system of the present invention, where compatibility is required for multiple standards, the system has... signals, including flag signals, are generated by each state machine to handle some of the processing within that state machine. Values carried in the standards can be used to access machine... its contents must be removed from the two wire interface to ensure that no further processing takes place using these 3 bytes. The decode register is emptied, and the value decode... 10. TOKENS

In the practice of the present invention, a token is a universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions and is adapted for use with a reconfigurable processing stage (RPS) which is a stage, which in response to a recognized token, reconfigures itself to perform various operations.

Tokens may be either position dependent or position independent upon the processing stages for performance of various functions. Tokens may also be metamorphic in that they can be altered by a processing stage and then passed down the pipeline for performance of further functions. Tokens may interact... other functions, and the specific interaction with a stage may be conditioned by the previous processing history of a stage.

A PICTURE(underscore)END token is a way of signalling the... through a fixed size, fixed width buffer.

The present invention is directed to a pipeline processing system which has a variable configuration which uses tokens and a two-wire system. The... do not use control tokens.

The control tokens are generated by circuitry within the decoder processor and emulate the operation of a number of different type standard-dependent signals passing into the serial pipeline processor for handling. The technique used is to study all the parameters of the multi-standards that are selected for processing by the serial processor and noting 1) their similarities, 2) their dissimilarities, 3) their needs and requirements and 4) selecting the correct token function to effectively process all of the standard signals sent into the serial processor. The functions of the tokens are to emulate the standards. A control token function is... the standard dependent signals and as an element to transmit control information through the pipeline processor.

In prior art system, a dedicated machine is designed according to well-known techniques to... tokens provide and make a sensible format for communicating information through the decompression circuit pipeline processor. In the design selected hereinafter and used in the preferred embodiment, each word of a... However, this is not a limitation on the invention, but on the magnitude of the processing steps elected to be accomplished by use of these tokens. It is to be noted... bit address for use in accessing the random access memories used throughout this serial decompression processor. This provides an additional degree of variability that facilitates a broad range of versatility.

As previously described, the DATA token carries data from one processing stage to the next. Consequently, the characteristics of this token change as it passes through... longest number of data bits because it needs to provide the most information to the

processing unit so that it can start the decompression with as much information as possible. Words which...to receive an address, it waits for the address generator to supply a valid address, processes that address and then sets the accept line high for one clock period. Thus, it... ..be read. This signal passes between two asynchronous clock regimes and, therefore, passes through three synchronizing flip flops.

Provided RAM2 312 is empty, the next item of data to arrive on...interesting.

In general, prediction data will be offset from the position of the block being processed as specified in the motion vectors in x and y. Thus, the block of data... ..address, 9. Data is read from this address and the x value is incremented. The process is repeated until the x value reaches its stop value, at which point, the y... ..is read, the x value is again incremented until it reaches its stop value. The process is repeated until both x and y values have reached their stop values. Thus, the...invention, is that additional information must be provided to the prediction filters to indicate what processing is required on the data. This consists of the following:

a "last byte" signal indicating... ..bit 0) is incremented and the x address (3 LSBS) is reset to zero. This process is repeated until 64 bytes have been read. With a 16 or 32 bit wide...register while its access register is set to zero, the results are undefined.

14. MICRO-PROCESSOR INTERFACE

A standard byte wide micro-processor interface (MPI) is used on all circuits with in the Spatial Decoder and Temporal Decoder... ..the parameter column. The actual specifications are shown in the respective columns min, max and units.

The DC operating conditions can be seen with reference to Table A.6.3. Here... ..the signal is present the maximum amount of time that this signal is available. The Units column gives the units of measurement used to describe the signals.

16. MPI WRITE TIMING

The general description of...a PICTURE(underscore)END token is decoded and forces the data in the coded data buffers to be applied to the Huffman decoder and video demultiplexor, the final picture can be... ..Consequently, the machine will not go into error recovery mode and will successfully continue to process the coded data.

A still further advantage of the use of a PICTURE(underscore)END token is that the serial pipeline processor will continue the processing of uninterrupted data. Through the use of a PICTURE(underscore)END token, the serial pipeline processor is configured to handle less than the expected amount of data and, therefore, continues processing. Typically, a prior art machine would stop itself because of an error condition. As previously... ..of the Huffman decode and Video Demultiplexor know the number of blocks that it will process during each picture recovery cycle. When the correct number of blocks do not arrive from... ..Each of the state machines recognizes a FLUSH control token as information not to be processed. Accordingly, the FLUSH token is used to fill up all of the remaining empty parts...less information than normally expected to decode the last picture. The Huffman decode circuit finishes processing the information contained in the last picture, and outputs this information through the DRAM interface... ..token, in accordance with the present invention, is used to pass through the entire pipeline processor and to ensure that the buffers are emptied and that other circuits are reconfigured to... ..underscore)END token, a padding word and a FLUSH token indicating to the serial pipeline processor that the picture processing for the current picture form is completed. Thereafter, the various state machines need reconfiguring to... ..FLUSH token resets each stage as it passes through, but-allows subsequent stages to continue processing. This prevents a loss of data. In other words, the FLUSH token is a variable... ..AFTER PICTURE

The STOP(underscore)AFTER(underscore)PICTURE function is employed to shut down the processing of the serial pipeline decompressing circuit at a logical point in its operation. At this... ..a picture, the STOP(underscore)AFTER(underscore)PICTURE operation signals the end of all current processing.

22. MULTI(underscore)STANDARD - SEARCH MODE

Another feature of the present invention is the use... ..MODE control token which is used to reconfigure the input to the serial pipeline processor to look at the incoming bit stream. When the search mode is set, the Start... ..combination of control tokens, and DATA tokens along with the reconfiguration circuits, to provide similar processing.

The use of search mode in the present invention is convenient in many situations including... ..video disc. In general, a search mode is convenient when the user interrupts the normal processing of the serial pipeline at a point where the machine does not expect such an...be the case.

In brief, the Huffman Decoder 321 works in conjunction with the other units shown in Figure 27. These other units are the Parser State Machine 322, the inshifter 323, the Index to Data unit 324, the ALU 325, and the Token Formatter 326. As described previously, connection between these blocks is governed by a two wire interface. A more detailed description of how these units function is subsequently described herein in greater detail, the focus here is on particular aspects... ..control certain functions of the Index to Data 324 and ALU 325. Control of these units by the Huffman Decoder is necessary for proper decoding of block-level information. Having the... ..further detail in the "More Detailed Description of the Invention" section.

The Index to Data unit 324 performs the second part of the multi-part algorithm. This unit contains a look up table that provides the actual Huffman decoded data. Entries in the...by detecting these in the Huffman Decoder 321, rather than in the Index to Data unit 324.

This index number is then passed to the Index to Data unit 324. In essence, the Index to Data unit is a look-up table. In accordance with one aspect of the algorithm, the look... ..format that JPEG specifies for transferring an alternate JPEG table.

From the Index to Data unit 324, the decoded index number or other data is passed, together with the accompanying control... ..the entering data to ensure that the DATA tokens are of the correct size for processing. In fact, the token stream can be corrected in some situations if the error is... ..an order that is useful for the decompression circuits, but not for the particular display unit being used. When a block of data enters the Buffer Manager, the Buffer Manager supplies...the output of the Spatial Decoder or Temporal Decoder and re-format it for a computer or display system. The details of this formatting will vary between applications. In a simple...Token. The DATA Token can have as many bits as are necessary for carrying out processing at a particular place in the system. All other Tokens ignore the extra bits.

A.3.2 The DATA Token

The DATA Token carries data from one processing stage to the next. Consequently, the characteristics of this Token change as it passes through... ..will be sufficient to collect DATA Tokens and to detect a few Tokens that provide synchronization information (such as PICTURE(underscore)START). In this regard, see subsequent sections A.16, "Connecting... ..from the data stream. This provides an alternative to doing the configuration via the micro processor interface.

A.3.4 Description of Tokens

This section documents the Tokens which are implemented... ..3.5.1. Note: JPEG requires a 2:1:1 structure for its macroblocks when processing 4:2:2 data. See Table A.3.5.

A.3.6 Special Token formats...either is low then the interface is taken to high impedance.

Note: on-chip data processing is not terminated when the DRAM interface is at high impedance. Therefore, errors will occur...decoded video's picture rate. Accordingly, this clock can be used to provide audio/video synchronization.

A.7.1 Spatial Decoder clock signals

The Spatial Decoder has two different (and potentially... ..in accordance with the present invention, must know what video standard is being input for processing. Thereafter, the system can accept either pre-existing Tokens or raw byte data which is...time a value is written into coded(underscore)data (7:0). Software is responsible for settling coded(underscore)extn to 0 before the last word of any Token is written to... ..0). The start of this new DATA Token then passes into the Spatial Decoder for processing.

Each time a new 8 bit value is written to coded(underscore)data (7:0... ..Detector analyses data in the DATA Tokens bit serially. The Detector's normal rate of processing is one bit per clock cycle (of coded(underscore)clock). Accordingly, it will typically decode a byte of coded data every 8 cycles of coded(underscore)clock. However, extra processing cycles are occasionally required, e.g., when a non-DATA Token is supplied or when main decoder(underscore)clock. Data transfer is synchronized to decoder(underscore)clock on-chip.

SECTION A.11 Start code detector

A.11.1... ..Code Detector. So, accessing these registers will be unreliable if the Start Code Detector is processing data. The user is responsible for ensuring that the Start Code Detector is halted before... ..Detector. In this case, the Tokens are passed through the Start Code Detector with no processing to other stages of the Spatial Decoder. These Tokens can only be inserted just before...result will be unpredictable if this is done when the Start Code Detector is actively processing data.

Discard all mode can be safely initiated after any of the Start Code Detector...start code non-alignment interrupt is suppressed.

In contrast, however, JPEG was designed for a computer environment where byte alignment is guaranteed. Therefore, marker codes should only be detected when byte... ..the other hand, was designed to meet the needs of both communications (bit serial) and

computer (byte oriented) systems. Start codes in MPEG data should normally be byte aligned. However, the...result will be unpredictable if this is done when the Start Code Detector is actively processing data. So, before initiating a start code search, the Start Code Detector should be stopped so no data is being processed. The Start Code Detector is always in this condition if any of the Start Code...the spatial video decoding circuits (inverse modeler, quantizer and DCT). This second logical buffer allows processing time to include a spread so as to accommodate processing pictures having varying amounts of data.

Both buffers are physically held in a single off the unit for all the above mentioned registers is a 512 bit block of data. Accordingly, the... ..until there is space in the buffer. If a buffer continues to be full, more processing stages "up steam" of the buffer will halt until the Spatial Decoder is unable to... ..converting coded data into Tokens started by the Start Code Detector. There are four main processing blocks in the Video Demux: Parser State Machine, Huffman decoder (including an ITOD), Macroblock counter... ..or state machine follows the syntax of the coded video data and instructs the other units. The Huffman decoder converts variable length coded (VLC) data into integers. The Macroblock counter keeps...

6/K/6 (Item 6 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Country	Number	Kind	Date		
Type		Pub. Date		Kind	Text
Available Text		Language		Update	Word Count
Total Word Count (Document A)					
Total Word Count (Document B)					
Total Word Count (All Documents)					

Specification: INTRODUCTION

The present invention is directed to improvements in methods and apparatus for decompression which operates to decompress and/or decode a plurality of differently encoded input... ..of the well known standards known as JPEG, MPEG and H.261.

A serial pipeline processing system of the present invention comprises a single two-wire bus used for carrying unique... ..to a plurality of adaptive decompression circuits and the like positioned as a reconfigurable pipeline processor.

PRIOR ART

One prior art system is described in United States Patent No. 5,216,724. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and a host processor. The device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus.

United States Patent No. 4,785... ..a known quad tree data structure.

United States Patent No. 5,122,875 discloses an apparatus for encoding/decoding an HDTV signal. The apparatus includes a compression circuit responsive to high definition video source signals for providing hierarchically layered... ..compressed video data of relatively greater and lesser importance to image reproduction respectively. A transport processor, responsive to the high and low priority codeword sequences, forms high and low priority transport... ..United States Patent No. 5,168,356 discloses a video signal encoding system that includes apparatus for segmenting encoded video data into transport blocks for signal transmission. ...in respective transport blocks.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening. This is accomplished by an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ..frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ..oscillators and the receiver can continuously receive each channel, then the receiver need not be synchronized with the transmitter. An FFT algorithm implements a fast discrete approximation to the continuous case in which the receiver synchronizes to the first frame and then acquires subsequent frames every frame period. The frame period... ..increasing the amount of data transmitted.

United States Patent No. 5,212,742 discloses an apparatus and method for processing video data for compression/decompression in real-time. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and host processor. Lastly, the device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus. The method handles assigning portions of the image for each of the processors to operate upon.

United States Patent No. 5,231,484 discloses a system and method... ..MPEG standards. Included are three cooperating components or subsystems that operate to variously adaptively pre-process the incoming digital motion video sequences, allocate bits to the pictures in a sequence, and... ..States Patent No. 5,267,334 discloses a method of removing frame redundancy in a computer system for a sequence of moving images. The method comprises detecting a first scene change... ..facing" keyframe or intraframe, and it is normally present in CCITT compressed video data. The process then comprises generating at least one intermediate compressed frame, the at least one intermediate compressed frame containing difference information from the first image for at least one image following...change, known as a "backward-facing" keyframe. The first keyframe and the at least one intermediate compressed frame are linked for forward play, and the second keyframe and the intermediate compressed frames are linked in reverse for reverse play. The intraframe may also be used... ..of complete scene information.

United States Patent No. 5,276,513 discloses a first circuit apparatus, comprising a given number of prior-art image-pyramid stages, together with a second circuit apparatus, comprising the same given number of novel motion-vector stages, perform cost-effective hierarchical motion analysis (HMA) in real-time, with minimum system processing delay and/or employing minimum system processing delay and/or employing minimum hardware structure. Specifically, the first and second circuit apparatus, in response to relatively high-resolution image data from an ongoing input series of successive... ..a relatively high frame rate (e.g., 30 frames per second), derives, after a certain processing-system delay, an ongoing output series of successive given pixel-density vector-data frames that... ..of successive image frames.

United States Patent No. 5,283,646 discloses a method and apparatus for enabling a real-time video encoding system to accurately deliver the desired number of... ..desired bit allocations.

The article, Chong, Yong M., A Data-Flow Architecture for Digital Image Processing, Wescon Technical Papers: No. 2 Oct./Nov. 1984, discloses a real-time signal processing system specifically designed for image processing. More particularly, a token based data-flow architecture is disclosed wherein the tokens are of... ..width having a fixed width address field. The system contains a plurality of identical flow processors connected in a ring fashion. The tokens contain a data field, a control field and a tag. The tag field of the token is further broken down into a processor address field and an identifier field. The processor address field is used to direct the tokens to the correct data-flow processor, and the identifier field is used to label the data such that the data-flow processor knows what to do with the data. In this way, the identifier field acts as an instruction for the data-flow processor. The system directs each token to a specific data-flow processor using a module number (MN). If the MN matches the MN of the particular stage... ..to locate the decoder in the preceding stage in order to pre-decode complex decoding processing and to alleviate critical path problems in the logic circuit. The elastic nature of the...of block signal in most cases.

United States Patent No. 4,903,018 discloses a process and data processing system for compressing and expanding structurally associated multiple data sequences. The process is particular to data sets in which an analysis is made of the structure in... ..data series on the basis of the order number of these data elements. The data processing system for performing the processes includes a storage matrix (26) and an index storage (28) having line addresses of the... ..the final actual video.

United States Patent No. 5,060,242 discloses an image signal processing system DPCM encodes the signal, then Huffman and run length encodes the signal to produce... ..tightly packed without gaps for efficient transmission

without loss of any data. The tightly packed apparatus has a barrel shifter with its shift modulus controlled by an accumulator receiving code word... ..OR gate is connected to the shifter, while a register is connected to the gate. Apparatus for processing a tightly packed and decorrelated digital signal has a barrel shifter and accumulator for unpacking... ..an inverse DCPM decoder.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening is accomplished by use of an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ..frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ..kernel matrix.

United States Patent No. 5,231,486 discloses a high definition video system processes a bitstream including high and low priority variable length coded Data words. The coded Data... ..packed High Priority Data and packed Low Priority Data by means of respective data packing units. The coded Data is continuously applied to both packing units. High Priority and Low Priority Length words indicating the bit lengths of high priority and... ..States Patent No. 5,287,178 discloses a video signal encoding system includes a signal processor for segmenting encoded video data into transport blocks having a header section and a packed data section. The system also includes reset control apparatus for releasing resets of system components, after a global system reset, in a prescribed non-simultaneous phased sequence to enable signal processing to commence in the prescribed sequence. The phased reset release sequence begins when valid data...United States Patent No. 5,142,380 to Sakagami et al. discloses an image compression apparatus

suitable for use with still images such as those formed by electronic still cameras using... ..and Q.

United States Patent No. 5,193,002 to Guichard et al. disclosed an apparatus for coding/decoding image signals in real time in conjunction with the CCITT standard H.261. A digital signal processor carries out direct quantization and reverse quantization.

United States Patent No. 5,241,383 to Chen et al. describes an apparatus with a pseudo-constant bit rate video coding achieved by an adjustable quantization parameter. The... ..relates to an improved pipeline system having an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along the pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the processing stages in the pipeline and interacting with selected stages in the pipeline for control data and/or combined control-data functions among the processing stages, so that the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing. In accordance with the invention, the processing stages may be configurable in response to recognition of at least one token. One of the processing stages may be a Start Code Detector which receives the input and generates and/or... ..and resetting the system, and a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards. The present invention... ..data and having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video pictures of varying data size can be controlled. Also in accordance with the invention, a processing stage receives the input data stream, the stage including means for recognizing specified bit stream patterns, whereby the processing stage facilitates random access and error recovery. The invention may also include a means for... ..invention also includes an inverse modeller stage, an inverse discrete cosine transform stage, and a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for processing data.

In addition, the present invention relates to an improved pipeline system having a Huffman...pipeline stage that incorporates a two-wire transfer control and also shows two consecutive pipeline processing stages with the two-wire transfer control;

Figures. 5a and 5b taken together depict one... ..shown in Figures. 8a and 8b.

Figure 10 is a block diagram of a reconfigurable processing stage;

Figure 11 is a block diagram of a spatial decoder;

Figure 12 is a... ..decoder including the prediction filters;

Figure 18 is a pictorial representation of the prediction filtering process;

Figure 19 shows a generalized representation of the macroblock structure;

Figure 20 shows a generalized... ..buffer;

Figure 25 is a pictorial diagram illustrating prediction data offset from the block being processed;

Figure 26 is a pictorial diagram illustrating prediction data offset by (1,1);

Figure 27...in general terms, the present invention provides an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along a pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the stages in the pipeline and interacting with selected stages in the pipeline for control, data and/or combined control-data functions among the processing stages, whereby the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing.

Each of the processing stages in the pipeline may include both primary and secondary storage, and the stages in processing stages for performance of functions or position independent of the processing stages for performance of functions.

In a pipeline machine, in accordance with the invention, the... ..altered by interfacing with the stages, and the tokens may interact with all of the processing stages in the pipeline or only with some but less than all of said processing stages. The tokens in the pipeline may interact with adjacent processing stages or with non-adjacent processing stages, and the tokens may reconfigure the processing stages. Such tokens may be position dependent for some functions and position independent for other... ..be Huffman coded.

In the improved pipeline machine, the tokens may be generated by a processing stage. Such pipeline tokens may include data for transfer to the processing stages or the tokens may be devoid of data. Some of the tokens may be identified as DATA tokens and provide data to the processing stages in the pipeline, while other tokens are identified as control tokens and only condition the processing stages in the pipeline, such conditioning including reconfiguring of the processing stages. Still other tokens may provide both data and conditioning to the processing stages in the pipeline. Some of said tokens may identify coding standards to the processing stages in the pipeline, whereas other tokens may operate independent of any coding standard among the processing stages. The tokens may be capable of successive alteration by the processing stages in the pipeline.

In accordance with the invention, the interactive flexibility of the tokens in cooperation with the processing stages facilitates greater functional diversity of the processing stages for resident structure in the pipeline, and the flexibility of the tokens facilitates system... ..or alteration. The tokens may be capable of facilitating a plurality of functions within any processing stage in the pipeline. Such pipeline tokens may be either hardware based or software based... ..system bandwidth in the pipeline. The tokens may provide data and control simultaneously to the processing stages in the pipeline.

The invention may include a pipeline processing machine for handling plurality of separately encoded bit streams arranged as a single serial bit... ..and for passing unrecognized control tokens along the pipeline, and a reconfigurable decode and parser processing means responsive to a recognized control token for reconfiguring a particular stage to handle an... ..be a pipeline system and the Start Code Detector may be positioned as the first processing stage in the pipeline.

The present invention also provides, in a system having a plurality of processing stages, a universal adaptation unit in the form of an interactive interfacing token for control and/or data functions among the processing stages, the token being a PICTURE(underscore)START code token for indicating that the start... ..The token may also be a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards.

The CODING(underscore... ..standard as JPEG, and/or any other appropriate picture standard. At least some of the processing stages reconfigure in response to the CODING(underscore)STANDARD token.

One of the processing stages in the system may be a Huffman decoder and parser and, upon receipt of... ..Data stage, and the parser stage may send an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... ..video data, having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video... ..controlled.

The system may include a spatial decoder having a two-wire interface interconnecting processing stages, the interface enabling serial processing for data and parallel processing for control.

As previously indicated, the system may further include a ROM having separate stored... ..of a plurality of picture standards, the programs being selectable by a token to facilitate processing for a plurality of different picture standards.

The spatial decoder system also includes a token... ..decoding stage and a parser stage for sending an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... ..The present invention also provides a pipeline system having an input data stream, and a processing stage for receiving the input data stream, the stage including means for recognizing specified bit... ..whereby said stage facilitates random access and error recovery. In accordance with the invention, the processing stage may be a start code detector and the bit stream patterns may include start... ..token and padding insures uniformity of word size. In accordance with the invention, a reconfigurable processing stage may be provided as a spatial decoder and the padding means adds to picture... ..that if the DATA token has less than the predetermined length, the padder circuit adds units of data to the DATA token until the predetermined length is achieved. A bypass circuit...1 tokens into a buffer, having a second predetermined width.

The invention also provides an apparatus for providing a time delay to a group of compressed pictures, the pictures corresponding to... ..and capable of delaying the words of data, is in communication with a control circuit intermediate the counter circuit and the inverse modeller circuit, the control circuit also communicating with the... ..inverse modeller stage and an inverse discrete cosine transform stage, the improvement characterized by a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for processing data.

In accordance with the invention, the token may be a QUANT(underscore)TABLE token for causing the processing stage to generate a quantization table.

The present invention also provides a Huffman decoder for... ..of bits used to represent an item of data.

DECODER: An embodiment of a decoding process.

DECODING (PROCESS): The process defined in this specification that reads an input coded bitstream and produces decoded pictures or... ..the same order in which they were presented at the input of the encoder.

ENCODING (PROCESS): A process, not specified in this specification, that reads a stream of input pictures or audio samples...to provide an estimate of the pel value or data element currently being decoded.

RECONFIGURABLE PROCESS STAGE (RPS): A stage, which in response to a recognized token, reconfigures itself to perform various operations.

SLICE: A series of macroblocks.

TOKEN: A universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions... ..indicates that the corresponding stage holds valid data, i.e., data that is to be processed in one of the pipeline stages. After processing (which may involve nothing more than a simple transfer without manipulation of the data) valid... ..present invention may be used with any number of pipeline stages. Furthermore, data may be processed in more than one stage and the processing time for different stages can differ.

In addition to clock and data signals (described below... ..other system. For example, the last pipeline stage may pass its data on to subsequent processing circuitry. The ACCEPT signal, which is illustrated as the lower of the two lines connecting... ..the minimum disturbance possible to other pipeline stages. Succeeding pipeline stages are allowed to continue processing and, therefore, this means that gaps open up in the stream of data following the... ..The data in the pipeline is encoded such that many different types of data are processed in the pipeline. This encoding accommodates data packets of variable size and the size of...the other hand, it may generate itself, all or part of the data to be processed in the pipeline. Indeed, as is explained below, a "stage" may contain arbitrary processing circuitry, including none at all (for simple passing of data) or entire systems (for example... ..values zero and 255 may not be used.

If such a picture were to be processed in a pipeline built in the practice of the present invention, then one of these... ..data must not be written over since it is data that must be saved for processing or use in a downstream device e.g., a pipeline stage, a device or a... ..connected to the pipeline upstream contains data D4 that is to be transferred into and processed in the pipeline. ...pipeline, in accordance with the preferred embodiments of the present invention, to "fill up" empty processing stages is highly advantageous since the processing stages in the pipeline thereby become decouple from one another. In other words, even though... ..data can be transferred into the pipeline and between stages even when one or more processing stages is blocked.

In the embodiment shown in Fig. 1, it is assumed that the...propagate all the way back to the beginning of the pipeline if there is some intermediate stage that is able to accept new data.

In the embodiment illustrated in Fig. 1...has been mentioned. It is to be further understood that each pipeline stage may also process the data it has received arbitrarily before passing it between its internal storage elements or... ..the portion of the pipeline that contains input and output storage elements and that arbitrarily processes data stored in its storage elements.

Furthermore, the "device" ...valid data, but also when a stage requires more than one clock phase to finish processing its data. This also can occur when it creates valid data in one or both... ..control the passage of data between adjacent storage elements. The VALID signal may also be processed in an analogous manner.

A great advantage of the two-wire interface (one wire for... ..In addition, two extra latches and a small number of gates are preferably added to process the ACCEPT and VALID signals that are associated with the data latches in

each half... application so requires. The interface in accordance with this embodiment can also be used to process analog signals.

As discussed previously, while other conventional timing arrangements may be used, the interface... circuit B1, which may be provided to convert output data from input latch LDIN into intermediate data, which is then later loaded in an output data latch LDOUT, which comprises the... is connected either directly as an input to the validation output latch LVOUT, or via intermediate logic devices or circuits that may alter the signal.

Similarly, the output validation signal QVOUT... to the input of the validation input latch QVIN of the following stage, or via intermediate devices or logic circuits, which may alter the validation signal. This ...word.

Preferred Data Structure - "tokens"

In the sample application shown in Fig. 4, each stage processes all input data, since there is no control circuitry that excludes any stage from allowing... are connected together in a relatively simple configuration. The simplest configuration is a pipeline of processing steps. For example, in the one shown in Fig. 1. The use of tokens, however... flows from left to right in the diagram. Data enters the machine and passes into processing Stage A. This may or may not modify the data and it then passes the... advantage of the tokens is their ability to achieve this kind of communication. Since any processing stage that does not recognize a token simply passes it on unaltered to the next... is transmitted along with the address and data fields in each token so that a processing stage can pass on a token (which can be of arbitrary length) without having to... be the first word of a new token.

Note that although the simple pipeline of processing stages is particularly useful, it will be appreciated that tokens may be applied to more complicated configurations of processing elements. An example of a more complicated processing element is described below.

It is not necessary, in accordance with the present invention, to... has extension bits. An example of this is a token that activates a stage that processes video quantization values stored in a quantization table (typically a memory device). For example, a... turn, is of great importance in video data pipeline systems since it ensures that all processing stages can be continuously running at full bandwidth.

In accordance with the present invention, in... some other chips in the set. This is advantageous both from the perspective of a customer and from that of a chip manufacturer. Even if modifications mean that all chips are... the end of a token (and hence the start of the next token) to be processed correctly (including simple non-manipulative transfer), even if the token is not recognized by the... block diagram of a pipeline stage whose function is as follows. If the stage is processing a predetermined token (known in this example as the DATA token), then it will duplicate... the address field of the DATA token. If, on the other hand, the stage is processing any other kind of token, it will delete every word. The overall effect is that... respective output signals:

In the duplication stage, the output from the data latch LDIN forms intermediate data referred to as MID(underscore)DATA. This intermediate data word is loaded into the data output latch LDOUT only when an intermediate acceptance signal (labeled "MID(underscore)ACCEPT" in Fig. 8a) is set HIGH.

The portion of... data. These include a "DATA(underscore)TOKEN" signal that indicates that the circuitry is currently processing a valid DATA Token, and a NOT(underscore)DUPLICATE signal which is used to control duplication of data. When the circuitry is processing a DATA Token, the NOT(underscore)DUPLICATE signal toggles between a HIGH and a LOW... the token to be duplicated once (but no more times). When the circuitry is not processing a valid DATA Token then the NOT(underscore)DUPLICATE signal is held in a HIGH state. Accordingly, this means that the token words that are being processed are not duplicated.

As Fig. 8a illustrates, the upper six bits of 8-bit intermediate data word and the output signal QI1 from the latch LI1 form inputs to a... explained further below.

Latch LO1 performs the function of latching the last value of the intermediate extension bit (labeled "MID(underscore)EXTN" and as signal S4), and it loads this value... ..and the DATA(underscore)TOKEN signal will become "0", indicating that the circuitry is not processing a DATA token.

If QI1 is "0" and SO is "0", thereby indicating a DATA... ..phase and the DATA(underscore)TOKEN signal will become "1", indicating that the circuitry is processing a DATA token.

The NOT(underscore)DUPLICATE signal (the output signal QO3) is similarly loaded...LVOUT at the same time that MID(underscore)DATA is loaded into LDOUT and the intermediate extension bit (signal S4) is loaded into LEOUT. Signal S5 is also combined with the...above. This has the effect that all tokens except the one that causes the duplication process will be deleted from the token stream, since a device connected to the output terminals (OUTDATA, OUTEXTN and OUTVALID) will not recognize these token words as valid data.

As before... ..and is duplicated.

Referring now more particularly to Figure 10, there is shown a reconfigurable process stage in accordance with one aspect of the present invention.

Input latches 34 receive an... ..the input latches 34 is passed as a first input over line 35 to a processing unit 36. A first output from the token decode subsystem 33 is passed over line 37 as a second input to the processing unit 36. A second output from the token decode 33 is passed over line 40 to an action identification unit 39. The action identification unit 39 also receives input from registers 43 and 44 over line 46. The registers 43... ..is determined by the history of tokens previously received. The output from the action identification unit 39 is passed over line 38 as a third input to the processing unit 36. The output from the processing unit 36 is passed to output latches 41. The output from the output latches 41 is... ..decoder 56 is passed over line 63 as an input to an Index to Data Unit (ITOD) 64. The Huffman decoder 56 and the ITOD 64 work together as a single logical unit. The output from the ITOD 64 is passed over line 65 to an arithmetic logic unit (ALU) 66. A first output from the ALU 66 is passed over line 67 to...blocks 133.

Referring to Figure 14b, in the JPEG and H.261 standards, the Common Intermediate Format (CIF) is used, wherein a picture 141 is encoded as 6 rows each containing... ..in a zigzag direction indicated by the arrow 144. The GOBs 142 are, in turn, processed row-by-row, left-to-right in each row.

Referring now to Figure 14c, it... ..in accordance with the practice of the present invention. A first picture 161 to be processed contains a first PICTURE(underscore)START token 162, first-picture information of indeterminate length 163, and a first PICTURE(underscore)END token 164. A second picture 165 to be processed contains a second PICTURE(underscore)START token 166, second picture information of indeterminate length 167... ..tokens 162 and 166 indicate the start of the pictures 161 and 165 to the processor. Likewise, the PICTURE(underscore)END tokens 164 and 168 signify the end of the pictures 161 and 165 to the processor. This allows the processor to process picture information 163 and 167 of variable lengths.

Referring to Figure 17, a split 171...Video Formatter (not shown in Figure 17).

Referring now to Figure 18, the prediction filtering process is illustrated. A forward picture 201 is passed over line 202 as a first input... ..the right of the value decode shift register 230, as indicated by area 231. This process eliminates overlapping start code images, as discussed below. A first output from the value decode... ..Code Detector. The Start Code Detector then receives a first data value image 244. Before processing the first data value image 244, the Start Code Detector may detect a second start... ..image 244 at a length 246. If this occurs, the Start Code Detector does not process the first data value image 244, and instead receives and processes a second data value image 247.

...line 1 of Table 600, whenever a "sequence start" image is received during H.261 processing or a "picture start" image is received during MPEG processing, the entire group of four control tokens is generated, each followed by its corresponding data...Picture Decoding

3. Motion Picture Decompression

4. RAM Memory Map

5. Bitstream Characteristics

6. Reconfigurable Processing Stage

7. Multi-Standard Coding

8. Multi-Standard Processing Circuit-2nd Mode of Operation

9. Start Code Detector

10. Tokens

11. DRAM Interface

12... ..described herein in greater detail) and reformatting this output for use, including display in a computer or other display systems, including a video display system. Implementation of this formatting varies significantly... ..the Spatial Decoder circuits.

The Spatial Decoder of the present invention performs all the required processing within a single picture. This reduces the redundancy within one picture.

The Temporal Decoder reduces... ..modeller 75, the inverse zig-zag 81 and the inverse DCT 83. The standard independent units within the Huffman decoder and parser include the ALU 66 and the token formatter 71.

Referring now to Figure 12, the standard-independent units include the DRAM interface 100, the fork 91, the FIFO register 96, the summer 98 and the output selector 106. The standard dependent units are the address generator 94, which is different in H.261 and in MPEG, and...much of the operation is very similar between the three different compression standards.

The next unit is the state machine 68 (Figure 11) located within the Huffman decoder and parser. Here... ..The same holds true for JPEG, which is a third,completely independent program.

The next unit is the Huffman decoder 56 which functions with the index to data unit 64. Those two units cooperate together to perform the Huffman decoding. Here, the algorithm that is used for Huffman... ..to the Huffman decoder at different times consistent with the standard in operation.

The last unit on the chip that is dependent on the compression standard is the inverse quantizer 79...an H.261 group of blocks and an MPEG slice. When H.261 data is processed after the Start Code Detector, each group of blocks is preceded by a slice(underscore... ..these standards have totally different sets of tables.

As previously indicated, most of the system units are compression standard independent. If a unit is standard independent, and such units need not remember what CODING(underscore)STANDARD is being processed. All of the units that are standard dependent remember the compression standard as the CODING(underscore)STANDARD token flows... ..CODING(underscore)STANDARD tokens at the Start Code Detector that is positioned as the first unit in the pipeline, this change of compression standard is readily handled. The token says a... ..found in the standard, i.e. from the bitstream into a prediction mode token. This processing is performed by the Huffman decoder and parser state machine, where it is easy to... ..to that token. By having these tokens and using them appropriately, the design of other units in the machine is simplified. Although there may be some complications in the program,

benefits... ..a first encoded signal (the MPEG or H.261 encoded video signal) in a pipeline processing system. The Temporal Decoder is not needed for JPEG decoding.

In this regard, the invention... ..the use of a single pipeline decoder and decompression system. The decoding and decompression pipeline processor is organized on a unique and special configuration which allows the handling of the multi... ..video signals through the use of techniques all compatible with the single pipeline decoder and processing system. The Spatial Decoder is combined with the Temporal Decoder, and the Video Formatter is...with only still pictures. The compression standard independent Spatial Decoder performs all of the data processing within the boundaries of a single picture. Such a decoder handles the spatial decompression of... ..to the multi-standard, configurable Video Formatter, which then provides an output to the display terminal. In a first sequence of similar pictures, each decompressed picture at the output of the... ..of control tokens and DATA tokens, in combination with a plurality of sequentially-positioned reconfigurable processing stages selected and organized to act as a standard-independent, reconfigurable-pipeline-processor.

With regard to JPEG decoding, a single Spatial Decoder with no off chip DRAM can... ..video. Accordingly, signals carried by DATA tokens pass directly through the Temporal Decoder without further processing when the Temporal Decoder is configured for a JPEG operation.

Another aspect of the present... ..for subsequent use in temporal decoding of subsequent pictures.

Generally, the Temporal Decoder performs the processing between pictures either earlier and/or later in time with reference to the picture currently... ..is distributed among several areas of DRAM in the sense that the decompressed output information, processed by the Spatial Decoder, is stored in other DRAM registers by other random access memories...first decoder circuit (the Spatial Decoder) directly to the Video Formatter for handling without signal processing delay.

The Temporal Decoder also reorders the blocks of picture data for display by a... ..from a selection of pictures which have arrived earlier or later than the picture under processing. When a picture is described in this context, it may mean any one of the... ..2. The result, i.e., the final decoded picture resulting from the addition of a process step performed by the decoder;

3. Previously decoded pictures read from the DRAM; and

4... ..START token and a subsequent PICTURE(underscore)END token.

After the picture data information is processed by the Temporal Decoder, it is either displayed or written back into a picture memory location. This information is then kept for further reference to be used in processing another different coded data picture.

Re-ordering of the MPEG encoded pictures for visual display...used to encode a referenced picture of a picture might be identified as being one unit long, another picture might be a number of units long, while still a third picture could be a fraction of that unit.

None of the existing standards (MPEG 1.2, JPEG, H.261) define a way of... ..picture rate, whereas the Video Formatter can handle a variable input picture rate.

6. RECONFIGURABLE PROCESSING STAGE

Referring again to Figure 10, the reconfigurable processing stage (RPS) comprises a token decode circuit 33 which is employed to receive the tokens... ..input latches 34. The output of the token decode circuit 33 is applied to a processing unit 36 over the two-wire interface 37 and an action identification circuit 39. The processing unit 36 is suitable for processing data under the control of the action identification circuit 39. After the processing is completed, the processing unit 36 connects such completed signals to the output, two-wire interface bus 40 through output... ..token decode circuit 33 are applied simultaneously to the action identification circuit 39 and the

processing unit 36. The action identification function as well as the RPS is described in further detail... ..not standard independent circuits. The data flows through the token decode circuit 33, through the

processing unit 36 and onto the two-wire interface circuit 42 through the output latches 41. If... ..wire interface 42 through the output circuit 41. The present invention operates as a pipeline processor having a two-wire interface for controlling the movement of control tokens through the pipeline... ..time, the token decode circuit 33 provides a proper flag or index signal to the processing unit 36 to alert it to the presence of the token being handled by the action identification circuit 39.

Control tokens may also be processed.

A more detailed description of the various types of tokens usable in the present invention...standard now passing through the state machine shown with reference to Figure 10.

Similarly, the processing unit 36 which is under the control of the action identification circuit 39 is now ready to process the information contained in the data fields of the DATA token when it is appropriate... ..action identification circuit 39 and is immediately followed by a DATA token which is then processed by the processing unit 36. The control token exits the output latches circuit 41 over the output two-wire interface 42 immediately preceding the DATA token which has been processed within the processing unit 36.

In the present invention, the action identification circuit, 39, is a state machine holding... ..show that the action can also be affected by the token that is currently being processed by the token decode circuit 33.

In general, there is shown token decoding and data processing in accordance with the present invention. The data processing is performed as configured by the action identification circuit 39. The action is affected by... ..information stored from previously decoded tokens in registers 43 and 44, the current token under processing, and the state and history information that the action identification unit 39 has itself acquired. A distinction is thereby shown between Control tokens and DATA tokens.

In any RPS, some tokens are viewed by that RPS unit as being Control tokens in that they affect the operation of the RPS presumably at... ..are viewed by the RPS as DATA tokens. Such DATA tokens contain information which is processed by the RPS in a way that is determined by the design of the particular... ..view of the same token. Some of the tokens might be viewed by one RPS unit as DATA Tokens while another RPS unit might decide that it is actually a Control Token. For example, the quantization table information... ..into a token called a quantization table token (QUANT(underscore)TABLE) which goes down the processing pipeline. As far as that machine is concerned, all of that was data; it was... ..sort of data into another sort of data, which is clearly a function of the processing performed by that portion of the machine. However, when that information gets to the inverse... ..present. This information is viewed as control information, and then that control information affects the processing that is done on subsequent DATA tokens because it affects the number that you multiply... ..important feature of the invention is that each of the stages of circuitry has the processing capability within it to be able to perform the necessary operations for each of the... ..operations are to be performed at a given time, come as tokens. There is one processing element that differs between the different stages to provide this capability. In the state machine...standard is and it looks up the parameters that it needs to apply to the processing elements in order to perform a proper operation. For example, the inverse quantizer will look... ..is set to 1 for a particular compression standard, and will apply that to its processing circuitry.

In a similar sense the Huffman decoder 56 has a number of tables within... ..MPEG video standard or the JPEG video standard. These three compression coding standards specify similar processes to be done on the arriving data, but the structure of the datastreams is different... ..token stream embodying the current coding standard. The control tokens are passed through the pipeline processor, and are used, i.e., decoded, in the state machines to which they are relevant... ..this regard, the DATA Tokens are treated in the same fashion, insofar as they are processed only in the

state machines that are configurable by the control tokens into processing such DATA Tokens. In the remaining state machines, they pass through unchanged.

More specifically, a... ..signals. The remaining portions of the token are used to indicate and identify the internal processing control function which is standard for all of the datastreams passing through the pipeline processor. In one form of the invention, the token extension is used to carry the current...accompanying data. As previously discussed, this information is utilized in the system to reconfigure the processing stage used to perform the function required by the various standards created for that purpose... ..picture number as indicated by the value.

The system also includes a multi-stage parallel processing pipeline operating under the principles of the two-wire interface previously described. Each of the... ..the token presently entering the state machine into the action identification circuit 39 or the processing unit 36, as appropriate. The processing unit has been previously reconfigured by the next previous control token into the form needed for handling the current coding standard, which is now entering the processing stage and carried by the next DATA token. Further, in accordance with this aspect of the invention, the succeeding state machines in the processing pipeline can be functioning under one coding standard, i.e., H.261, while a previous... ..tokens required to decode a number of coding standards with a fixed number of reconfigurable processing stages. More specifically, the PICTURE(underscore)END control token is employed because it is important... ..standard machine, it is necessary to create additional control tokens within the multi-standard pipeline processing machine which will then indicate which one of the standard decoding techniques to use. Such... ..and to push the current picture through the decoder to the display.

8. MULTI-STANDARD PROCESSING CIRCUIT - SECOND MODE OF OPERATION

A compression standard-dependent circuit, in the form of the...of the Start Code Detector will subsequently be discussed in further detail, as will the process of starting up of the decoder.

The aforementioned description has been concerned primarily ...the data which immediately follows according to the standard. However, in the multi-standard pipeline processing system of the present invention, where compatibility is required for multiple standards, the system has... ..signals, including flag signals, are generated by each state machine to handle some of the processing within that state machine. Values carried in the standards can be used to access machine... ..its contents must be removed from the two wire interface to ensure that no further processing takes place using these 3 bytes. The decode register is emptied, and the value decode... ..10. TOKENS

In the practice of the present invention, a token is a universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions and is adapted for use with a reconfigurable processing stage (RPS) which is a stage, which in response to a recognized token, reconfigures itself to perform various operations.

Tokens may be either position dependent or position independent upon the processing stages for performance of various functions. Tokens may also be metamorphic in that they can be altered by a processing stage and then passed down the pipeline for performance of further functions. Tokens may interact... ..other functions, and the specific interaction with a stage may be conditioned by the previous processing history of a stage.

A PICTURE(underscore)END token is a way of signalling the... ..through a fixed size, fixed width buffer.

The present invention is directed to a pipeline processing system which has a variable configuration which uses tokens and a two-wire system. The... ..do not use control tokens.

The control tokens are generated by circuitry within the decoder processor and emulate the operation of a number of different type standard-dependent signals passing into the serial pipeline processor for handling. The technique used is to study all the parameters of the multi-standards that are selected for processing by the serial processor and noting 1) their similarities, 2) their dissimilarities, 3) their needs and requirements and 4) selecting the correct token function to effectively process all of the standard signals sent into the serial processor. The functions of the tokens

are to emulate the standards. A control token function is... ..the standard dependent signals and as an element to transmit control information through the pipeline processor.

In prior art system, a dedicated machine is designed according to well-known techniques to... ..tokens provide and make a sensible format for communicating information through the decompression circuit pipeline processor. In the design selected hereinafter and used in the preferred embodiment, each word of a... ..However, this is not a limitation on the invention, but on the magnitude of the processing steps elected to be accomplished by use of these tokens. It is to be noted... ..bit address for use in accessing the random access memories used throughout this serial decompression processor. This provides an additional degree of variability that facilitates a broad range of versatility.

As previously described, the DATA token carries data from one processing stage to the next. Consequently, the characteristics of this token change as it passes through... ..longest number of data bits because it needs to provide the most information to the

processing unit so that it can start the decompression with as much information as possible. Words which...to receive an address, it waits for the address generator to supply a valid address, processes that address and then sets the accept line high for one clock period. Thus, it... ..be read. This signal passes between two asynchronous clock regimes and, therefore, passes through three synchronizing flip flops.

Provided RAM2 312 is empty, the next item of data to arrive on...interesting.

In general, prediction data will be offset from the position of the block being processed as specified in the motion vectors in x and y. Thus, the block of data... ..address, 9. Data is read from this address and the x value is incremented. The process is repeated until the x value reaches its stop value, at which point, the y... ..is read, the x value is again incremented until it reaches its stop value. The process is repeated until both x and y values have reached their stop values. Thus, the...invention, is that additional information must be provided to the prediction filters to indicate what processing is required on the data. This consists of the following:

a "last byte" signal indicating... ..bit 0) is incremented and the x address (3 LSBS) is reset to zero. This process is repeated until 64 bytes have been read. With a 16 or 32 bit wide...register while its access register is set to zero, the results are undefined.

14. MICRO-PROCESSOR INTERFACE

A standard byte wide micro-processor interface (MPI) is used on all circuits with in the Spatial Decoder and Temporal Decoder... ..the parameter column. The actual specifications are shown in the respective columns min, max and units.

The DC operating conditions can be seen with reference to Table A.6.3. Here... ..the signal is present the maximum amount of time that this signal is available. The Units column gives the units of measurement used to describe the signals.

16. MPI WRITE TIMING

The general description of...a PICTURE(underscore)END token is decoded and forces the data in the coded data buffers to be applied to the Huffman decoder and video demultiplexor, the final picture can be... ..Consequently, the machine will not go into error recovery mode and will successfully continue to process the coded data.

A still further advantage of the use of a PICTURE(underscore)END token is that the serial pipeline processor will continue the processing of uninterrupted data. Through the use of a PICTURE(underscore)END token, the serial pipeline processor is configured to handle less than the expected amount of data and, therefore, continues processing. Typically, a prior art machine would stop itself because of an error condition. As previously... ..of the

Huffman decode and Video Demultiplexor know the number of blocks that it will process during each picture recovery cycle. When the correct number of blocks do not arrive from... Each of the state machines recognizes a FLUSH control token as information not to be processed. Accordingly, the FLUSH token is used to fill up all of the remaining empty parts... less information than normally expected to decode the last picture. The Huffman decode circuit finishes processing the information contained in the last picture, and outputs this information through the DRAM interface... token, in accordance with the present invention, is used to pass through the entire pipeline processor and to ensure that the buffers are emptied and that other circuits are reconfigured to... underscore)END token, a padding word and a FLUSH token indicating to the serial pipeline processor that the picture processing for the current picture form is completed. Thereafter, the various state machines need reconfiguring to... FLUSH token resets each stage as it passes through, but- allows subsequent stages to continue processing. This prevents a loss of data. In other words, the FLUSH token is a variable... AFTER PICTURE

The STOP(underscore)AFTER(underscore)PICTURE function is employed to shut down the processing of the serial pipeline decompressing circuit at a logical point in its operation. At this... a picture, the STOP(underscore)AFTER(underscore)PICTURE operation signals the end of all current processing.

22. MULTI(underscore)STANDARD - SEARCH MODE

Another feature of the present invention is the use... underscore)MODE control token which is used to reconfigure the input to the serial pipeline processor to look at the incoming bit stream. When the search mode is set, the Start... combination of control tokens, and DATA tokens along with the reconfiguration circuits, to provide similar processing.

The use of search mode in the present invention is convenient in many situations including... video disc. In general, a search mode is convenient when the user interrupts the normal processing of the serial pipeline at a point where the machine does not expect such an... be the case.

In brief, the Huffman Decoder 321 works in conjunction with the other units shown in Figure 27. These other units are the Parser State Machine 322, the inshifter 323, the Index to Data unit 324, the ALU 325, and the Token Formatter 326. As described previously, connection between these blocks is governed by a two wire interface. A more detailed description of how these units function is subsequently described herein in greater detail, the focus here is on particular aspects... control certain functions of the Index to Data 324 and ALU 325. Control of these units by the Huffman Decoder is necessary for proper decoding of block-level information. Having the... further detail in the "More Detailed Description of the Invention" section.

The Index to Data unit 324 performs the second part of the multi-part algorithm. This unit contains a look up table that provides the actual Huffman decoded data. Entries in the... by detecting these in the Huffman Decoder 321, rather than in the Index to Data unit 324.

This index number is then passed to the Index to Data unit 324. In essence, the Index to Data unit is a look-up table. In accordance with one aspect of the algorithm, the look... format that JPEG specifies for transferring an alternate JPEG table.

From the Index to Data unit 324, the decoded index number or other data is passed, together with the accompanying control... the entering data to ensure that the DATA tokens are of the correct size for processing. In fact, the token stream can be corrected in some situations if the error is... an order that is useful for the decompression circuits, but not for the particular display unit being used. When a block of data enters the Buffer Manager, the Buffer Manager supplies... the output of the Spatial Decoder or Temporal Decoder and re-format it for a computer or display system. The details of this formatting will vary between applications. In a simple... Token. The DATA Token can have as many bits as are necessary for carrying out processing at a particular place in the system. All other Tokens ignore the extra bits.

A.3.2 The DATA Token

The DATA Token carries data from one processing stage to the next. Consequently, the characteristics of this Token change as it passes through... will be sufficient to collect DATA Tokens and to detect a few Tokens that provide synchronization information (such as PICTURE(underscore)START). In this regard, see subsequent sections A.16, "Connecting... from the data stream. This provides an alternative to doing the configuration via the micro processor interface.

A.3.4 Description of Tokens

This section documents the Tokens which are implemented... 3.5.1. Note: JPEG requires a 2:1:1 structure for its macroblocks when processing 4:2:2 data. See Table A.3.5.

A.3.6 Special Token formats...either is low then the interface is taken to high impedance.

Note: on-chip data processing is not terminated when the DRAM interface is at high impedance. Therefore, errors will occur...decoded video's picture rate. Accordingly, this clock can be used to provide audio/video synchronization.

A.7.1 Spatial Decoder clock signals

The Spatial Decoder has two different (and potentially... in accordance with the present invention, must know what video standard is being input for processing. Thereafter, the system can accept either pre-existing Tokens or raw byte data which is...time a value is written into coded(underscore)data (7:0). Software is responsible for settling coded(underscore)extn to 0 before the last word of any Token is written to... 0). The start of this new DATA Token then passes into the Spatial Decoder for processing.

Each time a new 8 bit value is written to coded(underscore)data (7:0... Detector analyses data in the DATA Tokens bit serially. The Detector's normal rate of processing is one bit per clock cycle (of coded(underscore)clock). Accordingly, it will typically decode a byte of coded data every 8 cycles of coded(underscore)clock. However, extra processing cycles are occasionally required, e.g., when a non-DATA Token is supplied or when the main decoder(underscore)clock. Data transfer is synchronized to decoder(underscore)clock on-chip.

SECTION A.11 Start code detector

A.11.1... Code Detector. So, accessing these registers will be unreliable if the Start Code Detector is processing data. The user is responsible for ensuring that the Start Code Detector is halted before... Detector. In this case, the Tokens are passed through the Start Code Detector with no processing to other stages of the Spatial Decoder. These Tokens can only be inserted just before...result will be unpredictable if this is done when the Start Code Detector is actively processing data.

Discard all mode can be safely initiated after any of the Start Code Detector...start code non-alignment interrupt is suppressed.

In contrast, however, JPEG was designed for a computer environment where byte alignment is guaranteed. Therefore, marker codes should only be detected when byte... the other hand, was designed to meet the needs of both communications (bit serial) and

computer (byte oriented) systems. Start codes in MPEG data should normally be byte aligned. However, the...result will be unpredictable if this is done when the Start Code Detector is actively processing data. So, before initiating a start code search, the Start Code Detector should be stopped so no data is being processed. The Start Code Detector is always in this condition if any of the Start Code...the spatial video decoding circuits (inverse modeler, quantizer and DCT). This second logical buffer allows processing time to include a spread so as to accommodate processing pictures having varying amounts of data.

Both buffers are physically held in a single off 1, the unit for all the above mentioned registers is a 512 bit block of data. Accordingly, the... ..until there is space in the buffer. If a buffer continues to be full, more processing stages "up steam" of the buffer will halt until the Spatial Decoder is unable to... ..converting coded data into Tokens started by the Start Code Detector. There are four main processing blocks in the Video Demux: Parser State Machine, Huffman decoder (including an ITOD), Macroblock counter... ..or state machine follows the syntax of the coded video data and instructs the other units. The Huffman decoder converts variable length coded (VLC) data into integers. The Macroblock counter keeps...

6/K/7 (Item 7 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Start code detecting apparatus for video data stream

Country	Number	Kind	Date
---------	--------	------	------

Abstract ...A2

an interactive interfacing token, defining a universal adaptation unit for control and/or data functions among said processing stages ; and

one of said stages receiving said input and adapted to generate and/or...

Type	Pub. Date	Kind	Text
Available Text	Language	Update	Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: INTRODUCTION

The present invention is directed to improvements in methods and apparatus for decompression which operates to decompress and/or decode a plurality of differently encoded input... ..of the well known standards known as JPEG, MPEG and H.261.

A serial pipeline processing system of the present invention comprises a single two-wire bus used for carrying unique... ..to a plurality of adaptive decompression circuits and the like positioned as a reconfigurable pipeline processor.

PRIOR ART

One prior art system is described in United States Patent No. 5,216,724. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and a host processor. The device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus.

United States Patent No. 4,785... ..a known quad tree data structure.

United States Patent No. 5,122,875 discloses an apparatus for encoding/decoding an HDTV signal. The apparatus includes a compression circuit responsive to high definition video source signals for providing hierarchically layered... compressed video data of relatively greater and lesser importance to image reproduction respectively. A transport processor, responsive to the high and low priority codeword sequences, forms high and low priority transport... United States Patent No. 5,168,356 discloses a video signal encoding system that includes apparatus for segmenting encoded video data into transport blocks for signal transmission. ...in respective transport blocks.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening. This is accomplished by an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... oscillators and the receiver can continuously receive each channel, then the receiver need not be synchronized with the transmitter. An FFT algorithm implements a fast discrete approximation to the continuous case in which the receiver synchronizes to the first frame and then acquires subsequent frames every frame period. The frame period... increasing the amount of data transmitted.

United States Patent No. 5,212,742 discloses an apparatus and method for processing video data for compression/decompression in real-time. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and host processor. Lastly, the device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus. The method handles assigning portions of the image for each of the processors to operate upon.

United States Patent No. 5,231,484 discloses a system and method... MPEG standards. Included are three cooperating components or subsystems that operate to variously adaptively pre-process the incoming digital motion video sequences, allocate bits to the pictures in a sequence, and... States Patent No. 5,267,334 discloses a method of removing frame redundancy in a computer system for a sequence of moving images. The method comprises detecting a first scene change... "forward-facing" keyframe or intraframe, and it is normally present in CCITT compressed video data. The process then comprises generating at least one intermediate compressed frame, the at least one intermediate compressed frame containing difference information from the first image for at least one image following... change, known as a "backward-facing" keyframe. The first keyframe and the at least one intermediate compressed frame are linked for forward play, and the second keyframe and the intermediate compressed frames are linked in reverse for reverse play. The intraframe may also be used... of complete scene information.

United States Patent No. 5,276,513 discloses a first circuit apparatus, comprising a given number of prior-art image-pyramid stages, together with a second circuit apparatus, comprising the same given number of novel motion-vector stages, perform cost-effective hierarchical motion analysis (HMA) in real-time, with minimum system processing delay and/or employing minimum system processing delay and/or employing minimum hardware structure. Specifically, the first and second circuit apparatus, in response to relatively high-resolution image data from an ongoing input series of successive... a relatively high frame rate (e.g., 30 frames per second), derives, after a certain processing-system delay, an ongoing output series of successive given pixel-density vector-data frames that... of successive image frames.

United States Patent No. 5,283,646 discloses a method and apparatus for enabling a real-time video encoding system to accurately deliver the desired number of... desired bit allocations.

The article, Chong, Yong M., A Data-Flow Architecture for Digital Image Processing, Wescon Technical Papers: No. 2 Oct./Nov. 1984, discloses a real-time signal processing system specifically designed for image processing.

More particularly, a token based data-flow architecture is disclosed wherein the tokens are of... width having a fixed width address field. The system contains a plurality of identical flow processors connected in a ring fashion. The tokens contain a data field, a control field and a tag. The tag field of the token is further broken down into a processor address field and an identifier field. The processor address field is used to direct the tokens to the correct data-flow processor, and the identifier field is used to label the data such that the data-flow processor knows what to do with the data. In this way, the identifier field acts as an instruction for the data-flow processor. The system directs each token to a specific data-flow processor using a module number (MN). If the MN matches the MN of the particular stage... to locate the decoder in the preceding stage in order to pre-decode complex decoding processing and to alleviate critical path problems in the logic circuit. The elastic nature of the... of block signal in most cases.

United States Patent No. 4,903,018 discloses a process and data processing system for compressing and expanding structurally associated multiple data sequences. The process is particular to data sets in which an analysis is made of the structure in... data series on the basis of the order number of these data elements. The data processing system for performing the processes includes a storage matrix (26) and an index storage (28) having line addresses of the... the final actual video.

United States Patent No. 5,060,242 discloses an image signal processing system DPCM encodes the signal, then Huffman and run length encodes the signal to produce... tightly packed without gaps for efficient transmission without loss of any data. The tightly packed apparatus has a barrel shifter with its shift modulus controlled by an accumulator receiving code word... OR gate is connected to the shifter, while a register is connected to the gate. Apparatus for processing a tightly packed and decorrelated digital signal has a barrel shifter and accumulator for unpacking... an inverse DCPM decoder.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening is accomplished by use of an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... kernel matrix.

United States Patent No. 5,231,486 discloses a high definition video system processes a bitstream including high and low priority variable length coded Data words. The coded Data... packed High Priority Data and packed Low Priority Data by means of respective data packing units. The coded Data is continuously applied to both packing units. High Priority and Low Priority Length words indicating the bit lengths of high priority and... States Patent No. 5,287,178 discloses a video signal encoding system includes a signal processor for segmenting encoded video data into transport blocks having a header section and a packed data section. The system also includes reset control apparatus for releasing resets of system components, after a global system reset, in a prescribed non-simultaneous phased sequence to enable signal processing to commence in the prescribed sequence. The phased reset release sequence begins when valid data... United States Patent No. 5,142,380 to Sakagami et al. discloses an image compression apparatus

suitable for use with still images such as those formed by electronic still cameras using... and Q.

United States Patent No. 5,193,002 to Guichard et al. disclosed an apparatus for coding/decoding image signals in real time in conjunction with the CCITT standard H.261. A digital signal processor carries out direct quantization and reverse quantization.

United States Patent No. 5,241,383 to Chen et al. describes an apparatus with a pseudo-constant bit rate video coding achieved by an adjustable quantization parameter. The... relates to an improved pipeline system having an

input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along the pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the processing stages in the pipeline and interacting with selected stages in the pipeline for control data and/or combined control-data functions among the processing stages, so that the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing. In accordance with the invention, the processing stages may be configurable in response to recognition of at least one token. One of the processing stages may be a Start Code Detector which receives the input and generates and/or... ..and resetting the system, and a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards. The present invention... ..data and having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video pictures of varying data size can be controlled. Also in accordance with the invention, a processing stage receives the input data stream, the stage including means for recognizing specified bit stream patterns, whereby the processing stage facilitates random access and error recovery. The invention may also include a means for... ..invention also includes an inverse modeller stage, an inverse discrete cosine transform stage, and a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for processing data.

In addition, the present invention relates to an improved pipeline system having a Huffman...pipeline stage that incorporates a two-wire transfer control and also shows two consecutive pipeline processing stages with the two-wire transfer control;

Figures. 5a and 5b taken together depict one... ..shown in Figures. 8a and 8b.

Figure 10 is a block diagram of a reconfigurable processing stage;

Figure 11 is a block diagram of a spatial decoder;

Figure 12 is a... ..decoder including the prediction filters;

Figure 18 is a pictorial representation of the prediction filtering process;

Figure 19 shows a generalized representation of the macroblock structure;

Figure 20 shows a generalized... ..buffer;

Figure 25 is a pictorial diagram illustrating prediction data offset from the block being processed;

Figure 26 is a pictorial diagram illustrating prediction data offset by (1,1);

Figure 27...in general terms, the present invention provides an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along a pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the stages in the pipeline and interacting with selected stages in the pipeline for control, data and/or combined control-data functions among the processing stages, whereby the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing.

Each of the processing stages in the pipeline may include both primary and secondary storage, and the stages in processing stages for performance of functions or position independent of the processing stages for performance of functions.

In a pipeline machine, in accordance with the invention, the... ..altered by interfacing with the stages, and the tokens may interact with all of the processing stages in the pipeline or only with some but less than all of said processing stages. The tokens in the pipeline may interact with adjacent processing stages or with non-adjacent processing

stages, and the tokens may reconfigure the processing stages. Such tokens may be position dependent for some functions and position independent for other... ..be Huffman coded.

In the improved pipeline machine, the tokens may be generated by a processing stage. Such pipeline tokens may include data for transfer to the processing stages or the tokens may be devoid of data. Some of the tokens may be identified as DATA tokens and provide data to the processing stages in the pipeline, while other tokens are identified as control tokens and only condition the processing stages in the pipeline, such conditioning including reconfiguring of the processing stages. Still other tokens may provide both data and conditioning to the processing stages in the pipeline. Some of said tokens may identify coding standards to the processing stages in the pipeline, whereas other tokens may operate independent of any coding standard among the processing stages. The tokens may be capable of successive alteration by the processing stages in the pipeline.

In accordance with the invention, the interactive flexibility of the tokens in cooperation with the processing stages facilitates greater functional diversity of the processing stages for resident structure in the pipeline, and the flexibility of the tokens facilitates system... ..or alteration. The tokens may be capable of facilitating a plurality of functions within any processing stage in the pipeline. Such pipeline tokens may be either hardware based or software based... ..system bandwidth in the pipeline. The tokens may provide data and control simultaneously to the processing stages in the pipeline.

The invention may include a pipeline processing machine for handling plurality of separately encoded bit streams arranged as a single serial bit... ..and for passing unrecognized control tokens along the pipeline, and a reconfigurable decode and parser processing means responsive to a recognized control token for reconfiguring a particular stage to handle an... ..be a pipeline system and the Start Code Detector may be positioned as the first processing stage in the pipeline.

The present invention also provides, in a system having a plurality of processing stages, a universal adaptation unit in the form of an interactive interfacing token for control and/or data functions among the processing stages, the token being a PICTURE(underscore)START code token for indicating that the start... ..The token may also be a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards.

The CODING(underscore... ..standard as JPEG, and/or any other appropriate picture standard. At least some of the processing stages reconfigure in response to the CODING(underscore)STANDARD token.

One of the processing stages in the system may be a Huffman decoder and parser and, upon receipt of... ..Data stage, and the parser stage may send an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... ..video data, having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video... ..controlled.

The system may include a spatial decoder having a two-wire interface interconnecting processing stages, the interface enabling serial processing for data and parallel processing for control.

As previously indicated, the system may further include a ROM having separate stored... ..of a plurality of picture standards, the programs being selectable by a token to facilitate processing for a plurality of different picture standards.

The spatial decoder system also includes a token... ..decoding stage and a parser stage for sending an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... ..The present invention also provides a pipeline system having an input data stream, and a processing stage for receiving the input data stream, the stage including means for recognizing specified bit... ..whereby said

stage facilitates random access and error recovery. In accordance with the invention, the processing stage may be a start code detector and the bit stream patterns may include start... token and padding insures uniformity of word size. In accordance with the invention, a reconfigurable processing stage may be provided as a spatial decoder and the padding means adds to picture... that if the DATA token has less than the predetermined length, the padder circuit adds units of data to the DATA token until the predetermined length is achieved. A bypass circuit...1 tokens into a buffer, having a second predetermined width.

The invention also provides an apparatus for providing a time delay to a group of compressed pictures, the pictures corresponding to... and capable of delaying the words of data, is in communication with a control circuit intermediate the counter circuit and the inverse modeller circuit, the control circuit also communicating with the... inverse modeller stage and an inverse discrete cosine transform stage, the improvement characterized by a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for processing data.

In accordance with the invention, the token may be a QUANT(underscore)TABLE token for causing the processing stage to generate a quantization table.

The present invention also provides a Huffman decoder for... of bits used to represent an item of data.

DECODER: An embodiment of a decoding process.

DECODING (PROCESS): The process defined in this specification that reads an input coded bitstream and produces decoded pictures or... the same order in which they were presented at the input of the encoder.

ENCODING (PROCESS): A process, not specified in this specification, that reads a stream of input pictures or audio samples...to provide an estimate of the pel value or data element currently being decoded.

RECONFIGURABLE PROCESS STAGE (RPS): A stage, which in response to a recognized token, reconfigures itself to perform various operations.

SLICE: A series of macroblocks.

TOKEN: A universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions... indicates that the corresponding stage holds valid data, i.e., data that is to be processed in one of the pipeline stages. After processing (which may involve nothing more than a simple transfer without manipulation of the data) valid... present invention may be used with any number of pipeline stages. Furthermore, data may be processed in more than one stage and the processing time for different stages can differ.

In addition to clock and data signals (described below... other system. For example, the last pipeline stage may pass its data on to subsequent processing circuitry. The ACCEPT signal, which is illustrated as the lower of the two lines connecting... the minimum disturbance possible to other pipeline stages. Succeeding pipeline stages are allowed to continue processing and, therefore, this means that gaps open up in the stream of data following the... The data in the pipeline is encoded such that many different types of data are processed in the pipeline. This encoding accommodates data packets of variable size and the size of...the other hand, it may generate itself, all or part of the data to be processed in the pipeline. Indeed, as is explained below, a "stage" may contain arbitrary processing circuitry, including none at all (for simple passing of data) or entire systems (for example... values zero and 255 may not be used.

If such a picture were to be processed in a pipeline built in the practice of the present invention, then one of these... data must not be written over since it is data that must be saved for processing or use in a downstream device e.g., a pipeline stage, a device or a... connected to the pipeline upstream contains data D4 that is to be transferred into and processed in the pipeline. ...pipeline, in accordance with the preferred embodiments of the present invention, to

"fill up" empty processing stages is highly advantageous since the processing stages in the pipeline thereby become decouple from one another. In other words, even though... data can be transferred into the pipeline and between stages even when one or more processing stages is blocked.

In the embodiment shown in Fig. 1, it is assumed that the...propagate all the way back to the beginning of the pipeline if there is some intermediate stage that is able to accept new data.

In the embodiment illustrated in Fig. 1...has been mentioned. It is to be further understood that each pipeline stage may also process the data it has received arbitrarily before passing it between its internal storage elements or... the portion of the pipeline that contains input and output storage elements and that arbitrarily processes data stored in its storage elements.

Furthermore, the "device" ...valid data, but also when a stage requires more than one clock phase to finish processing its data. This also can occur when it creates valid data in one or both... control the passage of data between adjacent storage elements. The VALID signal may also be processed in an analogous manner.

A great advantage of the two-wire interface (one wire for... In addition, two extra latches and a small number of gates are preferably added to process the ACCEPT and VALID signals that are associated with the data latches in each half... application so requires. The interface in accordance with this embodiment can also be used to process analog signals.

As discussed previously, while other conventional timing arrangements may be used, the interface... circuit B1, which may be provided to convert output data from input latch LDIN into intermediate data, which is then later loaded in an output data latch LDOUT, which comprises the... is connected either directly as an input to the validation output latch LVOUT, or via intermediate logic devices or circuits that may alter the signal.

Similarly, the output validation signal QVOUT... to the input of the validation input latch QVIN of the following stage, or via intermediate devices or logic circuits, which may alter the validation signal. This ...word.

Preferred Data Structure - "tokens"

In the sample application shown in Fig. 4, each stage processes all input data, since there is no control circuitry that excludes any stage from allowing... are connected together in a relatively simple configuration. The simplest configuration is a pipeline of processing steps. For example, in the one shown in Fig. 1. The use of tokens, however...flows from left to right in the diagram. Data enters the machine and passes into processing Stage A. This may or may not modify the data and it then passes the... advantage of the tokens is their ability to achieve this kind of communication. Since any processing stage that does not recognize a token simply passes it on unaltered to the next... is transmitted along with the address and data fields in each token so that a processing stage can pass on a token (which can be of arbitrary length) without having to... be the first word of a new token.

Note that although the simple pipeline of processing stages is particularly useful, it will be appreciated that tokens may be applied to more complicated configurations of processing elements. An example of a more complicated processing element is described below.

It is not necessary, in accordance with the present invention, to... has extension bits. An example of this is a token that activates a stage that processes video quantization values stored in a quantization table (typically a memory device). For example, a...turn, is of great importance in video data pipeline systems since it ensures that all processing stages can be continuously running at full bandwidth.

In accordance to the present invention, in... some other chips in the set. This is advantageous both from the perspective of a customer and from that of a chip manufacturer. Even if modifications mean that all chips are...the end of a token (and hence the start of the next token) to be processed correctly (including simple non-manipulative transfer), even if the token is not recognized by the... block diagram of a pipeline stage whose function is as follows. If the stage is processing a predetermined token (known in this example as the DATA token), then it will

duplicate... ..the address field of the DATA token. If, on the other hand, the stage is processing any other kind of token, it will delete every word. The overall effect is that... ..respective output signals:

In the duplication stage, the output from the data latch LDIN forms intermediate data referred to as MID(underscore)DATA. This intermediate data word is loaded into the data output latch LDOUT only when an intermediate acceptance signal (labeled "MID(underscore)ACCEPT" in Fig. 8a) is set HIGH.

The portion of... ..data. These include a "DATA(underscore)TOKEN" signal that indicates that the circuitry is currently processing a valid DATA Token, and a NOT(underscore)DUPLICATE signal which is used to control duplication of data. When the circuitry is processing a DATA Token, the NOT(underscore)DUPLICATE signal toggles between a HIGH and a LOW... ..the token to be duplicated once (but no more times). When the circuitry is not processing a valid DATA Token then the NOT(underscore)DUPLICATE signal is held in a HIGH state. Accordingly, this means that the token words that are being processed are not duplicated.

As Fig. 8a illustrates, the upper six bits of 8-bit intermediate data word and the output signal QI1 from the latch LI1 form inputs to a... ..explained further below.

Latch LO1 performs the function of latching the last value of the intermediate extension bit (labeled "MID(underscore)EXTN" and as signal S4), and it loads this value... ..and the DATA(underscore)TOKEN signal will become "0", indicating that the circuitry is not processing a DATA token.

If QI1 is "0" and SO is "0", thereby indicating a DATA... ..phase and the DATA(underscore)TOKEN signal will become "1", indicating that the circuitry is processing a DATA token.

The NOT(underscore)DUPLICATE signal (the output signal QO3) is similarly loaded...LVOUT at the same time that MID(underscore)DATA is loaded into LDOUT and the intermediate extension bit (signal S4) is loaded into LEOUT. Signal S5 is also combined with the...above. This has the effect that all tokens except the one that causes the duplication process will be deleted from the token stream, since a device connected to the output terminals (OUTDATA, OUTEXTN and OUTVALID) will not recognize these token words as valid data.

As before... ..and is duplicated.

Referring now more particularly to Figure 10, there is shown a reconfigurable process stage in accordance with one aspect of the present invention.

Input latches 34 receive an... ..the input latches 34 is passed as a first input over line 35 to a processing unit 36. A first output from the token decode subsystem 33 is passed over line 37 as a second input to the processing unit 36. A second output from the token decode 33 is passed over line 40 to an action identification unit 39. The action identification unit 39 also receives input from registers 43 and 44 over line 46. The registers 43... ..is determined by the history of tokens previously received. The output from the action identification unit 39 is passed over line 38 as a third input to the processing unit 36. The output from the processing unit 36 is passed to output latches 41. The output from the output latches 41 is... ..decoder 56 is passed over line 63 as an input to an Index to Data Unit (ITOD) 64. The Huffman decoder 56 and the ITOD 64 work together as a single logical unit. The output from the ITOD 64 is passed over line 65 to an arithmetic logic unit (ALU) 66. A first output from the ALU 66 is passed over line 67 to...blocks 133.

Referring to Figure 14b, in the JPEG and H.261 standards, the Common Intermediate Format (CIF) is used, wherein a picture 141 is encoded as 6 rows each containing... ..in a zigzag direction indicated by the arrow 144. The GOBs 142 are, in turn, processed row-by-row, left-to-right in each row.

Referring now to Figure 14c, it... ..in accordance with the practice of the present invention. A first picture 161 to be processed contains a first PICTURE(underscore)START token 162, first-picture information of indeterminate length 163, and a first PICTURE(underscore)END token 164. A second picture 165 to be processed contains a second PICTURE(underscore)START token 166, second picture information of indeterminate length 167... ..tokens 162 and 166 indicate the start of the pictures 161 and 165 to the processor. Likewise, the PICTURE(underscore)END tokens 164 and 168 signify the end of the pictures 161 and 165 to the processor. This allows the processor to process picture information 163 and 167 of variable lengths.

Referring to Figure 17, a split 171...Video Formatter (not shown in Figure 17).

Referring now to Figure 18, the prediction filtering process is illustrated. A forward picture 201 is passed over line 202 as a first input... ..the right of the value decode shift register 230, as indicated by area 231. This process eliminates overlapping start code images, as discussed below. A first output from the value decode... ..Code Detector. The Start Code Detector then receives a first data value image 244. Before processing the first data value image 244, the Start Code Detector may detect a second start... ..image 244 at a length 246. If this occurs, the Start Code Detector does not process the first data value image 244, and instead receives and processes a second data value image 247.

...line 1 of Table 600, whenever a "sequence start" image is received during H.261 processing or a "picture start" image is received during MPEG processing, the entire group of four control tokens is generated, each followed by its corresponding data...Picture Decoding

3. Motion Picture Decompression

4. RAM Memory Map

5. Bitstream Characteristics

6. Reconfigurable Processing Stage

7. Multi-Standard Coding

8. Multi-Standard Processing Circuit-2nd Mode of Operation

9. Start Code Detector

10. Tokens

11. DRAM Interface

12... ..described herein in greater detail) and reformatting this output for use, including display in a computer or other display systems, including a video display system. Implementation of this formatting varies significantly... ..the Spatial Decoder circuits.

The Spatial Decoder of the present invention performs all the required processing within a single picture. This reduces the redundancy within one picture.

The Temporal Decoder reduces... ..modeller 75, the inverse zig-zag 81 and the inverse DCT 83. The standard independent units within the Huffman decoder and parser include the ALU 66 and the token formatter 71.

Referring now to Figure 12, the standard-independent units include the DRAM interface 100, the fork 91, the FIFO register 96, the summer 98 and the output selector 106. The standard dependent units are the address generator 94, which is different in H.261 and in MPEG, and...much of the operation is very similar between the three different compression standards.

The next unit is the state machine 68 (Figure 11) located within the Huffman decoder and parser. Here... ..The same holds true for JPEG, which is a third,completely independent program.

The next unit is the Huffman decoder 56 which functions with the index to data unit 64. Those two units cooperate together to perform the Huffman decoding. Here, the algorithm that is used for Huffman... ..to the Huffman decoder at different times consistent with the standard in operation.

The last unit on the chip that is dependent on the compression standard is the inverse quantizer 79...an H.261 group of blocks and an MPEG slice. When H.261 data is processed after the Start Code Detector, each group of blocks is preceded by a slice(underscore... ..these standards have totally different sets of tables.

As previously indicated, most of the system units are compression standard independent. If a unit is standard independent, and such units need not remember what CODING(underscore)STANDARD is being processed. All of the units that are standard dependent remember the compression standard as the CODING(underscore)STANDARD token flows... ..CODING(underscore)STANDARD tokens at the Start Code Detector that is positioned as the first unit in the pipeline, this change of compression standard is readily handled. The token says a... ..found in the standard, i.e. from the bitstream into a prediction mode token. This processing is performed by the Huffman decoder and parser state machine, where it is easy to... ..to that token. By having these tokens and using them appropriately, the design of other units in the machine is simplified. Although there may be some complications in the program, benefits... ..a first encoded signal (the MPEG or H.261 encoded video signal) in a pipeline processing system. The Temporal Decoder is not needed for JPEG decoding.

In this regard, the invention... ..the use of a single pipeline decoder and decompression system. The decoding and decompression pipeline processor is organized on a unique and special configuration which allows the handling of the multi... ..video signals through the use of techniques all compatible with the single pipeline decoder and processing system. The Spatial Decoder is combined with the Temporal Decoder, and the Video Formatter is...with only still pictures. The compression standard independent Spatial Decoder performs all of the data processing within the boundaries of a single picture. Such a decoder handles the spatial decompression of... ..to the multi-standard, configurable Video Formatter, which then provides an output to the display terminal. In a first sequence of similar pictures, each decompressed picture at the output of the... ..of control tokens and DATA tokens, in combination with a plurality of sequentially-positioned reconfigurable processing stages selected and organized to act as a standard-independent, reconfigurable-pipeline-processor.

With regard to JPEG decoding, a single Spatial Decoder with no off chip DRAM can... ..video. Accordingly, signals carried by DATA tokens pass directly through the Temporal Decoder without further processing when the Temporal Decoder is configured for a JPEG operation.

Another aspect of the present... ..for subsequent use in temporal decoding of subsequent pictures.

Generally, the Temporal Decoder performs the processing between pictures either earlier and/or later in time with reference to the picture currently... ..is distributed among several areas of DRAM in the sense that the decompressed output information, processed by the Spatial Decoder, is stored in other DRAM registers by other random access memories...first decoder circuit (the Spatial Decoder) directly to the Video Formatter for handling without signal processing delay.

The Temporal Decoder also reorders the blocks of picture data for display by a... ..from a selection of pictures which have arrived earlier or later than the picture under processing. When a picture is described in this context, it may mean any one of the... ..2. The result, i.e., the final decoded picture resulting from the addition of a process step performed by the decoder;

3. Previously decoded pictures read from the DRAM; and

4... ..START token and a subsequent PICTURE(underscore)END token.

After the picture data information is processed by the Temporal Decoder, it is either displayed or written back into a picture memory location. This information is then kept for further reference to be used in processing another different coded data picture.

Re-ordering of the MPEG encoded pictures for visual display...used to encode a referenced picture of a picture might be identified as being one unit long, another picture might be a number of units long, while still a third picture could be a fraction of that unit.

None of the existing standards (MPEG 1.2, JPEG, H.261) define a way of... picture rate, whereas the Video Formatter can handle a variable input picture rate.

6. RECONFIGURABLE PROCESSING STAGE

Referring again to Figure 10, the reconfigurable processing stage (RPS) comprises a token decode circuit 33 which is employed to receive the tokens... input latches 34. The output of the token decode circuit 33 is applied to a processing unit 36 over the two-wire interface 37 and an action identification circuit 39. The processing unit 36 is suitable for processing data under the control of the action identification circuit 39. After the processing is completed, the processing unit 36 connects such completed signals to the output, two-wire interface bus 40 through output... token decode circuit 33 are applied simultaneously to the action identification circuit 39 and the processing unit 36. The action identification function as well as the RPS is described in further detail... not standard independent circuits. The data flows through the token decode circuit 33, through the

processing unit 36 and onto the two-wire interface circuit 42 through the output latches 41. If... wire interface 42 through the output circuit 41. The present invention operates as a pipeline processor having a two-wire interface for controlling the movement of control tokens through the pipeline... time, the token decode circuit 33 provides a proper flag or index signal to the processing unit 36 to alert it to the presence of the token being handled by the action identification circuit 39.

Control tokens may also be processed.

A more detailed description of the various types of tokens usable in the present invention... standard now passing through the state machine shown with reference to Figure 10.

Similarly, the processing unit 36 which is under the control of the action identification circuit 39 is now ready to process the information contained in the data fields of the DATA token when it is appropriate... action identification circuit 39 and is immediately followed by a DATA token which is then processed by the processing unit 36. The control token exits the output latches circuit 41 over the output two-wire interface 42 immediately preceding the DATA token which has been processed within the processing unit 36.

In the present invention, the action identification circuit, 39, is a state machine holding... show that the action can also be affected by the token that is currently being processed by the token decode circuit 33.

In general, there is shown token decoding and data processing in accordance with the present invention. The data processing is performed as configured by the action identification circuit 39. The action is affected by... information stored from previously decoded tokens in registers 43 and 44, the current token under processing, and the state and history information that the action identification unit 39 has itself acquired. A distinction is thereby shown between Control tokens and DATA tokens.

In any RPS, some tokens are viewed by that RPS unit as being Control tokens in that they affect the operation of the RPS presumably at... are viewed by the RPS as DATA tokens. Such DATA tokens contain information which is processed by the RPS in a way that is determined by the design of the particular... view of the same token. Some of the tokens might be viewed by one RPS unit as DATA Tokens while another RPS unit might decide that it is actually a Control Token. For example, the quantization table information... into a token called a quantization table token (QUANT(underscore)TABLE) which goes down the processing pipeline. As far as that machine is concerned,

all of that was data; it was... ..sort of data into another sort of data, which is clearly a function of the processing performed by that portion of the machine. However, when that information gets to the inverse... ..present. This information is viewed as control information, and then that control information affects the processing that is done on subsequent DATA tokens because it affects the number that you multiply... ..important feature of the invention is that each of the stages of circuitry has the processing capability within it to be able to perform the necessary operations for each of the... ..operations are to be performed at a given time, come as tokens. There is one processing element that differs between the different stages to provide this capability. In the state machine...standard is and it looks up the parameters that it needs to apply to the processing elements in order to perform a proper operation. For example, the inverse quantizer will look... ..is set to 1 for a particular compression standard, and will apply that to its processing circuitry.

In a similar sense the Huffman decoder 56 has a number of tables within... ..MPEG video standard or the JPEG video standard. These three compression coding standards specify similar processes to be done on the arriving data, but the structure of the datastreams is different... ..token stream embodying the current coding standard. The control tokens are passed through the pipeline processor, and are used, i.e., decoded, in the state machines to which they are relevant... ..this regard, the DATA Tokens are treated in the same fashion, insofar as they are processed only in the state machines that are configurable by the control tokens into processing such DATA Tokens. In the remaining state machines, they pass through unchanged.

More specifically, a... ..signals. The remaining portions of the token are used to indicate and identify the internal processing control function which is standard for all of the datastreams passing through the pipeline processor. In one form of the invention, the token extension is used to carry the current...accompanying data. As previously discussed, this information is utilized in the system to reconfigure the processing stage used to perform the function required by the various standards created for that purpose... ..picture number as indicated by the value.

The system also includes a multi-stage parallel processing pipeline operating under the principles of the two-wire interface previously described. Each of the... ..the token presently entering the state machine into the action identification circuit 39 or the processing unit 36, as appropriate. The processing unit has been previously reconfigured by the next previous control token into the form needed for handling the current coding standard, which is now entering the processing stage and carried by the next DATA token. Further, in accordance with this aspect of the invention, the succeeding state machines in the processing pipeline can be functioning under one coding standard, i.e., H.261, while a previous... ..tokens required to decode a number of coding standards with a fixed number of reconfigurable processing stages. More specifically, the PICTURE(underscore)END control token is employed because it is important... ..standard machine, it is necessary to create additional control tokens within the multi-standard pipeline processing machine which will then indicate which one of the standard decoding techniques to use. Such... ..and to push the current picture through the decoder to the display.

8. MULTI-STANDARD PROCESSING CIRCUIT - SECOND MODE OF OPERATION

A compression standard-dependent circuit, in the form of the...of the Start Code Detector will subsequently be discussed in further detail, as will the process of starting up of the decoder.

The aforementioned description has been concerned primarily ...the data which immediately follows according to the standard. However, in the multi-standard pipeline processing system of the present invention, where compatibility is required for multiple standards, the system has... ..signals, including flag signals, are generated by each state machine to handle some of the processing within that state machine. Values carried in the standards can be used to access machine... ..its contents must be removed from the two wire interface to ensure that no further processing takes place using these 3 bytes. The decode register is emptied, and the value decode... ..10. TOKENS

In the practice of the present invention, a token is a universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions and is adapted for use with a reconfigurable processing stage (RPS) which is a stage, which in response to a recognized token, reconfigures itself to perform various operations.

Tokens may be either position dependent or position independent upon the processing stages for performance of various functions. Tokens may also be metamorphic in that they can be altered by a processing stage and then passed down the pipeline for performance of further functions. Tokens may interact... other functions, and the specific interaction with a stage may be conditioned by the previous processing history of a stage.

A PICTURE(underscore)END token is a way of signalling the... through a fixed size, fixed width buffer.

The present invention is directed to a pipeline processing system which has a variable configuration which uses tokens and a two-wire system. The... do not use control tokens.

The control tokens are generated by circuitry within the decoder processor and emulate the operation of a number of different type standard-dependent signals passing into the serial pipeline processor for handling. The technique used is to study all the parameters of the multi-standards that are selected for processing by the serial processor and noting 1) their similarities, 2) their dissimilarities, 3) their needs and requirements and 4) selecting the correct token function to effectively process all of the standard signals sent into the serial processor. The functions of the tokens are to emulate the standards. A control token function is... the standard dependent signals and as an element to transmit control information through the pipeline processor.

In prior art system, a dedicated machine is designed according to well-known techniques to... tokens provide and make a sensible format for communicating information through the decompression circuit pipeline processor. In the design selected hereinafter and used in the preferred embodiment, each word of a... However, this is not a limitation on the invention, but on the magnitude of the processing steps elected to be accomplished by use of these tokens. It is to be noted... bit address for use in accessing the random access memories used throughout this serial decompression processor. This provides an additional degree of variability that facilitates a broad range of versatility.

As previously described, the DATA token carries data from one processing stage to the next. Consequently, the characteristics of this token change as it passes through... longest number of data bits because it needs to provide the most information to the

processing unit so that it can start the decompression with as much information as possible. Words which...to receive an address, it waits for the address generator to supply a valid address, processes that address and then sets the accept line high for one clock period. Thus, it... be read. This signal passes between two asynchronous clock regimes and, therefore, passes through three synchronizing flip flops.

Provided RAM2 312 is empty, the next item of data to arrive on...interesting.

In general, prediction data will be offset from the position of the block being processed as specified in the motion vectors in x and y. Thus, the block of data... address, 9. Data is read from this address and the x value is incremented. The process is repeated until the x value reaches its stop value, at which point, the y... is read, the x value is again incremented until it reaches its stop value. The process is repeated until both x and y values have reached their stop values. Thus, the...invention, is that additional information must be provided to the prediction filters to indicate what processing is required on the data. This consists of the following:

a "last byte" signal indicating... bit 0) is incremented and the x address (3 LSBS) is reset to zero. This process is repeated until 64 bytes have been read. With a 16 or 32 bit wide...register while its access register is set to zero, the results are undefined.

14. MICRO-PROCESSOR INTERFACE

A standard byte wide micro-processor interface (MPI) is used on all circuits with in the Spatial Decoder and Temporal Decoder... the parameter column. The actual specifications are shown in the respective columns min, max and units.

The DC operating conditions can be seen with reference to Table A.6.3. Here... the signal is present the maximum amount of time that this signal is available. The Units column gives the units of measurement used to describe the signals.

16. MPI WRITE TIMING

The general description of...a PICTURE(underscore)END token is decoded and forces the data in the coded data buffers to be applied to the Huffman decoder and video demultiplexor, the final picture can be... Consequently, the machine will not go into error recovery mode and will successfully continue to process the coded data.

A still further advantage of the use of a PICTURE(underscore)END token is that the serial pipeline processor will continue the processing of uninterrupted data. Through the use of a PICTURE(underscore)END token, the serial pipeline processor is configured to handle less than the expected amount of data and, therefore, continues processing. Typically, a prior art machine would stop itself because of an error condition. As previously... of the Huffman decode and Video Demultiplexor know the number of blocks that it will process during each picture recovery cycle. When the correct number of blocks do not arrive from... Each of the state machines recognizes a FLUSH control token as information not to be processed. Accordingly, the FLUSH token is used to fill up all of the remaining empty parts... Huffman Decoder and Video Demultiplexor. In this way, the FLUSH token is like padding for buffers.

The Token Decoder in the Huffman circuit recognizes the FLUSH token and ignores the pseudo... less information than normally expected to decode the last picture. The Huffman decode circuit finishes processing the information contained in the last picture, and outputs this information through the DRAM interface... token, in accordance with the present invention, is used to pass through the entire pipeline processor and to ensure that the buffers are emptied and that other circuits are reconfigured to... (underscore)END token, a padding word and a FLUSH token indicating to the serial pipeline processor that the picture processing for the current picture form is completed. Thereafter, the various state machines need reconfiguring to... FLUSH token resets each stage as it passes through, but-allows subsequent stages to continue processing. This prevents a loss of data. In other words, the FLUSH token is a variable... AFTER PICTURE

The STOP(underscore)AFTER(underscore)PICTURE function is employed to shut down the processing of the serial pipeline decompressing circuit at a logical point in its operation. At this... a picture, the STOP(underscore)AFTER(underscore)PICTURE operation signals the end of all current processing.

22. MULTI(underscore)STANDARD - SEARCH MODE

Another feature of the present invention is the use... (underscore)MODE control token which is used to reconfigure the input to the serial pipeline processor to look at the incoming bit stream. When the search mode is set, the Start... combination of control tokens, and DATA tokens along with the reconfiguration circuits, to provide similar processing.

The use of search mode in the present invention is convenient in many situations including... video disc. In general, a search mode is convenient when the user interrupts the normal processing of the serial pipeline at a point where the machine does not expect such an...be the case.

In brief, the Huffman Decoder 321 works in conjunction with the other units shown in Figure 27. These other units are the Parser State Machine 322, the inshifter 323, the Index to Data unit 324, the ALU 325, and the Token Formatter 326. As described previously, connection between these blocks is governed by a two wire interface. A more detailed description of how these units function is subsequently described herein in greater detail, the focus

here is on particular aspects... ..control certain functions of the Index to Data 324 and ALU 325. Control of these units by the Huffman Decoder is necessary for proper decoding of block-level information. Having the... ..further detail in the "More Detailed Description of the Invention" section.

The Index to Data unit 324 performs the second part of the multi-part algorithm. This unit contains a look up table that provides the actual Huffman decoded data. Entries in the...by detecting these in the Huffman Decoder 321, rather than in the Index to Data unit 324.

This index number is then passed to the Index to Data unit 324. In essence, the Index to Data unit is a look-up table. In accordance with one aspect of the algorithm, the look... ..format that JPEG specifies for transferring an alternate JPEG table.

From the Index to Data unit 324, the decoded index number or other data is passed, together with the accompanying control... ..the entering data to ensure that the DATA tokens are of the correct size for processing. In fact, the token stream can be corrected in some situations if the error is... ..an order that is useful for the decompression circuits, but not for the particular display unit being used. When a block of data enters the Buffer Manager, the Buffer Manager supplies...the output of the Spatial Decoder or Temporal Decoder and re-format it for a computer or display system. The details of this formatting will vary between applications. In a simple...Token. The DATA Token can have as many bits as are necessary for carrying out processing at a particular place in the system. All other Tokens ignore the extra bits.

A.3.2 The DATA Token

The DATA Token carries data from one processing stage to the next. Consequently, the characteristics of this Token change as it passes through... ..will be sufficient to collect DATA Tokens and to detect a few Tokens that provide synchronization information (such as PICTURE(underscore)START). In this regard, see subsequent sections A.16, "Connecting... ..from the data stream. This provides an alternative to doing the configuration via the micro processor interface.

A.3.4 Description of Tokens

This section documents the Tokens which are implemented... ..3.5.1. Note: JPEG requires a 2:1:1 structure for its macroblocks when processing 4:2:2 data. See Table A.3.5.

A.3.6 Special Token formats...either is low then the interface is taken to high impedance.

Note: on-chip data processing is not terminated when the DRAM interface is at high impedance. Therefore, errors will occur...decoded video's picture rate. Accordingly, this clock can be used to provide audio/video synchronization.

A.7.1 Spatial Decoder clock signals

The Spatial Decoder has two different (and potentially... ..in accordance with the present invention, must know what video standard is being input for processing. Thereafter, the system can accept either pre-existing Tokens or raw byte data which is...time a value is written into coded(underscore)data (7:0). Software is responsible for settling coded(underscore)extn to 0 before the last word of any Token is written to... ..0). The start of this new DATA Token then passes into the Spatial Decoder for processing.

Each time a new 8 bit value is written to coded(underscore)data (7:0... ..Detector analyses data in the DATA Tokens bit serially. The Detector's normal rate of processing is one bit per clock cycle (of coded(underscore)clock). Accordingly, it will typically decode a byte of coded data every 8 cycles of coded(underscore)clock. However, extra processing cycles are occasionally required, e.g., when a non-DATA Token is supplied or when the main decoder(underscore)clock. Data transfer is synchronized to decoder(underscore)clock on-chip.

SECTION A.11 Start code detector

A.11.1... ..Code Detector. So, accessing these registers will be unreliable if the Start Code Detector is processing data. The user is responsible for ensuring that the Start Code Detector is halted before... ..Detector. In this case, the Tokens are passed through the Start Code Detector with no processing to other stages of the Spatial Decoder. These Tokens can only be inserted just before...result will be unpredictable if this is done when the Start Code Detector is actively processing data.

Discard all mode can be safely initiated after any of the Start Code Detector...start code non-alignment interrupt is suppressed.

In contrast, however, JPEG was designed for a computer environment where byte alignment is guaranteed. Therefore, marker codes should only be detected when byte... ..the other hand, was designed to meet the needs of both communications (bit serial) and computer (byte oriented) systems. Start codes in MPEG data should normally be byte aligned. However, the...result will be unpredictable if this is done when the Start Code Detector is actively processing data. So, before initiating a start code search, the Start Code Detector should be stopped so no data is being processed. The Start Code Detector is always in this condition if any of the Start Code...the spatial video decoding circuits (inverse modeler, quantizer and DCT). This second logical buffer allows processing time to include a spread so as to accommodate processing pictures having varying amounts of data.

Both buffers are physically held in a single off 1.1, the unit for all the above mentioned registers is a 512 bit block of data. Accordingly, the... ..until there is space in the buffer. If a buffer continues to be full, more processing stages "up steam" of the buffer will halt until the Spatial Decoder is unable to... ..converting coded data into Tokens started by the Start Code Detector. There are four main processing blocks in the Video Demux: Parser State Machine, Huffman decoder (including an ITOD), Macroblock counter... ..or state machine follows the syntax of the coded video data and instructs the other units. The Huffman decoder converts variable length coded (VLC) data into integers. The Macroblock counter keeps...

Claims:

1. In a system having an input and an output and a plurality of processing stages between the input and the output, the improvement comprising:

an interactive interfacing token, defining a universal adaptation unit for control and/or data functions among said processing stages ; and

one of said stages receiving said input and adapted to generate and/or... ..a token.

11. The system as recited in claim 10, and further comprising a reconfigurable processing stage as a spatial decoder; and said means for padding adds to data being handled...

6/K/8 (Item 8 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Start code detecting apparatus for video data stream

Country	Number	Kind	Date
---------	--------	------	------

Abstract ...A2

a control circuit intermediate and in communication with said counter circuit and said inverse modeller circuit ;
said counter circuit...

Type	Pub. Date	Kind	Text
Available Text	Language	Update	Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: INTRODUCTION

The present invention is directed to improvements in methods and apparatus for decompression which operates to decompress and/or decode a plurality of differently encoded input... ..of the well known standards known as JPEG, MPEG and H.261.

A serial pipeline processing system of the present invention comprises a single two-wire bus used for carrying unique... ..to a plurality of adaptive decompression circuits and the like positioned as a reconfigurable pipeline processor.

PRIOR ART

One prior art system is described in United States Patent No. 5,216,724. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and a host processor. The device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus.

United States Patent No. 4,785... ..a known quad tree data structure.

United States Patent No. 5,122,875 discloses an apparatus for encoding/decoding an HDTV signal. The apparatus includes a compression circuit responsive to high definition video source signals for providing hierarchically layered... ..compressed video data of relatively greater and lesser importance to image reproduction respectively. A transport processor, responsive to the high and low priority codeword sequences, forms high and low priority transport... ..United States Patent No. 5,168,356 discloses a video signal encoding system that includes apparatus for segmenting encoded video data into transport blocks for signal transmission. ...in respective transport blocks.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening. This is accomplished by an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ..frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ..oscillators and the receiver can continuously receive each channel, then the receiver need not be synchronized with the transmitter. An FFT algorithm implements a fast discrete approximation to the continuous case in which the receiver synchronizes to the first frame and then acquires subsequent frames every frame period. The frame period... ..increasing the amount of data transmitted.

United States Patent No. 5,212,742 discloses an apparatus and method for processing video data for compression/decompression in real-time. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and

host processor. Lastly, the device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus. The method handles assigning portions of the image for each of the processors to operate upon.

United States Patent No. 5,231,484 discloses a system and method... MPEG standards. Included are three cooperating components or subsystems that operate to variously adaptively pre-process the incoming digital motion video sequences, allocate bits to the pictures in a sequence, and... States Patent No. 5,267,334 discloses a method of removing frame redundancy in a computer system for a sequence of moving images. The method comprises detecting a first scene change... "backward-facing" keyframe or intraframe, and it is normally present in CCITT compressed video data. The process then comprises generating at least one intermediate compressed frame, the at least one intermediate compressed frame containing difference information from the first image for at least one image following...change, known as a "backward-facing" keyframe. The first keyframe and the at least one intermediate compressed frame are linked for forward play, and the second keyframe and the intermediate compressed frames are linked in reverse for reverse play. The intraframe may also be used... of complete scene information.

United States Patent No. 5,276,513 discloses a first circuit apparatus, comprising a given number of prior-art image-pyramid stages, together with a second circuit apparatus, comprising the same given number of novel motion-vector stages, perform cost-effective hierarchical motion analysis (HMA) in real-time, with minimum system processing delay and/or employing minimum system processing delay and/or employing minimum hardware structure. Specifically, the first and second circuit apparatus, in response to relatively high-resolution image data from an ongoing input series of successive... a relatively high frame rate (e.g., 30 frames per second), derives, after a certain processing-system delay, an ongoing output series of successive given pixel-density vector-data frames that... of successive image frames.

United States Patent No. 5,283,646 discloses a method and apparatus for enabling a real-time video encoding system to accurately deliver the desired number of... desired bit allocations.

The article, Chong, Yong M., A Data-Flow Architecture for Digital Image Processing, Wescon Technical Papers: No. 2 Oct./Nov. 1984, discloses a real-time signal processing system specifically designed for image processing. More particularly, a token based data-flow architecture is disclosed wherein the tokens are of... width having a fixed width address field. The system contains a plurality of identical flow processors connected in a ring fashion. The tokens contain a data field, a control field and a tag. The tag field of the token is further broken down into a processor address field and an identifier field. The processor address field is used to direct the tokens to the correct data-flow processor, and the identifier field is used to label the data such that the data-flow processor knows what to do with the data. In this way, the identifier field acts as an instruction for the data-flow processor. The system directs each token to a specific data-flow processor using a module number (MN). If the MN matches the MN of the particular stage... to locate the decoder in the preceding stage in order to pre-decode complex decoding processing and to alleviate critical path problems in the logic circuit. The elastic nature of the... of block signal in most cases.

United States Patent No. 4,903,018 discloses a process and data processing system for compressing and expanding structurally associated multiple data sequences. The process is particular to data sets in which an analysis is made of the structure in... data series on the basis of the order number of these data elements. The data processing system for performing the processes includes a storage matrix (26) and an index storage (28) having line addresses of the... the final actual video.

United States Patent No. 5,060,242 discloses an image signal processing system DPCM encodes the signal, then Huffman and run length encodes the signal to produce... tightly packed without gaps for efficient transmission without loss of any data. The tightly packed apparatus has a barrel shifter with its shift modulus controlled by an accumulator receiving code word... OR gate is connected to the shifter, while a register is connected to the gate. Apparatus for processing a tightly packed and decorrelated digital signal has a barrel shifter and accumulator for unpacking... an inverse DCPM decoder.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening is accomplished by use of an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ..frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ..kernel matrix.

United States Patent No. 5,231,486 discloses a high definition video system processes a bitstream including high and low priority variable length coded Data words. The coded Data... ..packed High Priority Data and packed Low Priority Data by means of respective data packing units. The coded Data is continuously applied to both packing units. High Priority and Low Priority Length words indicating the bit lengths of high priority and... ..States Patent No. 5,287,178 discloses a video signal encoding system includes a signal processor for segmenting encoded video data into transport blocks having a header section and a packed data section. The system also includes reset control apparatus for releasing resets of system components, after a global system reset, in a prescribed non-simultaneous phased sequence to enable signal processing to commence in the prescribed sequence. The phased reset release sequence begins when valid data...United States Patent No. 5,142,380 to Sakagami et al. discloses an image compression apparatus

suitable for use with still images such as those formed by electronic still cameras using... ..and Q.

United States Patent No. 5,193,002 to Guichard et al. disclosed an apparatus for coding/decoding image signals in real time in conjunction with the CCITT standard H.261. A digital signal processor carries out direct quantization and reverse quantization.

United States Patent No. 5,241,383 to Chen et al. describes an apparatus with a pseudo-constant bit rate video coding achieved by an adjustable quantization parameter. The... ..relates to an improved pipeline system having an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along the pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the processing stages in the pipeline and interacting with selected stages in the pipeline for control data and/or combined control-data functions among the processing stages, so that the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing. In accordance with the invention, the processing stages may be configurable in response to recognition of at least one token. One of the processing stages may be a Start Code Detector which receives the input and generates and/or... ..and resetting the system, and a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards. The present invention... ..data and having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video pictures of varying data size can be controlled. Also in accordance with the invention, a processing stage receives the input data stream, the stage including means for recognizing specified bit stream patterns, whereby the processing stage facilitates random access and error recovery. The invention may also include a means for... ..invention also includes an inverse modeller stage, an inverse discrete cosine transform stage, and a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for processing data.

In addition, the present invention relates to an improved pipeline system having a Huffman...pipeline stage that incorporates a two-wire transfer control and also shows two consecutive pipeline processing stages with the two-wire transfer control;

Figures. 5a and 5b taken together depict one... ..shown in Figures. 8a and 8b.

Figure 10 is a block diagram of a reconfigurable processing stage;

Figure 11 is a block diagram of a spatial decoder;

Figure 12 is a... ..decoder including the prediction filters;

Figure 18 is a pictorial representation of the prediction filtering process;

Figure 19 shows a generalized representation of the macroblock structure;

Figure 20 shows a generalized... ..buffer;

Figure 25 is a pictorial diagram illustrating prediction data offset from the block being processed;

Figure 26 is a pictorial diagram illustrating prediction data offset by (1,1);

Figure 27...in general terms, the present invention provides an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along a pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the stages in the pipeline and interacting with selected stages in the pipeline for control, data and/or combined control-data functions among the processing stages, whereby the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing.

Each of the processing stages in the pipeline may include both primary and secondary storage, and the stages in processing stages for performance of functions or position independent of the processing stages for performance of functions.

In a pipeline machine, in accordance with the invention, the... ..altered by interfacing with the stages, and the tokens may interact with all of the processing stages in the pipeline or only with some but less than all of said processing stages. The tokens in the pipeline may interact with adjacent processing stages or with non-adjacent processing stages, and the tokens may reconfigure the processing stages. Such tokens may be position dependent for some functions and position independent for other... ..be Huffman coded.

In the improved pipeline machine, the tokens may be generated by a processing stage. Such pipeline tokens may include data for transfer to the processing stages or the tokens may be devoid of data. Some of the tokens may be identified as DATA tokens and provide data to the processing stages in the pipeline, while other tokens are identified as control tokens and only condition the processing stages in the pipeline, such conditioning including reconfiguring of the processing stages. Still other tokens may provide both data and conditioning to the processing stages in the pipeline. Some of said tokens may identify coding standards to the processing stages in the pipeline, whereas other tokens may operate independent of any coding standard among the processing stages. The tokens may be capable of successive alteration by the processing stages in the pipeline.

In accordance with the invention, the interactive flexibility of the tokens in cooperation with the processing stages facilitates greater functional diversity of the processing stages for resident structure in the pipeline, and the flexibility of the tokens facilitates system... ..or alteration. The tokens may be capable of facilitating a plurality of functions within any processing stage in the pipeline. Such pipeline tokens may be either hardware based or software based... ..system bandwidth in the pipeline. The tokens may provide data and control simultaneously to the processing stages in the pipeline.

The invention may include a pipeline processing machine for handling plurality of separately encoded bit streams arranged as a single serial bit... ..and for passing unrecognized control tokens along the pipeline, and a reconfigurable decode and parser processing means responsive to a recognized control token for reconfiguring a particular stage to handle an... ..be a pipeline system and the Start Code Detector may be positioned as the first processing stage in the pipeline.

The present invention also provides, in a system having a plurality of processing stages, a universal adaptation unit in the form of an interactive interfacing token for control and/or data functions among the processing stages, the token being a PICTURE(underscore)START code token for indicating that the start... The token may also be a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards.

The CODING(underscore)... standard as JPEG, and/or any other appropriate picture standard. At least some of the processing stages reconfigure in response to the CODING(underscore)STANDARD token.

One of the processing stages in the system may be a Huffman decoder and parser and, upon receipt of... Data stage, and the parser stage may send an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... video data, having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video... controlled.

The system may include a spatial decoder having a two-wire interface interconnecting processing stages, the interface enabling serial processing for data and parallel processing for control.

As previously indicated, the system may further include a ROM having separate stored... of a plurality of picture standards, the programs being selectable by a token to facilitate processing for a plurality of different picture standards.

The spatial decoder system also includes a token... decoding stage and a parser stage for sending an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... The present invention also provides a pipeline system having an input data stream, and a processing stage for receiving the input data stream, the stage including means for recognizing specified bit... whereby said stage facilitates random access and error recovery. In accordance with the invention, the processing stage may be a start code detector and the bit stream patterns may include start... token and padding insures uniformity of word size. In accordance with the invention, a reconfigurable processing stage may be provided as a spatial decoder and the padding means adds to picture... that if the DATA token has less than the predetermined length, the padder circuit adds units of data to the DATA token until the predetermined length is achieved. A bypass circuit... 1 tokens into a buffer, having a second predetermined width.

The invention also provides an apparatus for providing a time delay to a group of compressed pictures, the pictures corresponding to... and capable of delaying the words of data, is in communication with a control circuit intermediate the counter circuit and the inverse modeller circuit, the control circuit also communicating with the... inverse modeller stage and an inverse discrete cosine transform stage, the improvement characterized by a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for

processing data.

In accordance with the invention, the token may be a QUANT(underscore)TABLE token for causing the processing stage to generate a quantization table.

The present invention also provides a Huffman decoder for... of bits used to represent an item of data.

DECODER: An embodiment of a decoding process.

DECODING (PROCESS): The process defined in this specification that reads an input coded bitstream and produces decoded pictures or... the same order in which they were presented at the input of the encoder.

ENCODING (PROCESS): A process, not specified in this specification, that reads a stream of input pictures or audio samples...to provide an estimate of the pel value or data element currently being decoded.

RECONFIGURABLE PROCESS STAGE (RPS): A stage, which in response to a recognized token, reconfigures itself to perform various operations.

SLICE: A series of macroblocks.

TOKEN: A universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions... ..indicates that the corresponding stage holds valid data, i.e., data that is to be processed in one of the pipeline stages. After processing (which may involve nothing more than a simple transfer without manipulation of the data) valid... ..present invention may be used with any number of pipeline stages. Furthermore, data may be processed in more than one stage and the processing time for different stages can differ.

In addition to clock and data signals (described below... ..other system. For example, the last pipeline stage may pass its data on to subsequent processing circuitry. The ACCEPT signal, which is illustrated as the lower of the two lines connecting... ..the minimum disturbance possible to other pipeline stages. Succeeding pipeline stages are allowed to continue processing and, therefore, this means that gaps open up in the stream of data following the... ..The data in the pipeline is encoded such that many different types of data are processed in the pipeline. This encoding accommodates data packets of variable size and the size of...the other hand, it may generate itself, all or part of the data to be processed in the pipeline. Indeed, as is explained below, a "stage" may contain arbitrary processing circuitry, including none at all (for simple passing of data) or entire systems (for example... ..values zero and 255 may not be used.

If such a picture were to be processed in a pipeline built in the practice of the present invention, then one of these... ..data must not be written over since it is data that must be saved for processing or use in a downstream device e.g., a pipeline stage, a device or a... ..connected to the pipeline upstream contains data D4 that is to be transferred into and processed in the pipeline. ...pipeline, in accordance with the preferred embodiments of the present invention, to "fill up" empty processing stages is highly advantageous since the processing stages in the pipeline thereby become decouple from one another. In other words, even though... ..data can be transferred into the pipeline and between stages even when one or more processing stages is blocked.

In the embodiment shown in Fig. 1, it is assumed that the...propagate all the way back to the beginning of the pipeline if there is some intermediate stage that is able to accept new data.

In the embodiment illustrated in Fig. 1...has been mentioned. It is to be further understood that each pipeline stage may also process the data it has received arbitrarily before passing it between its internal storage elements or... ..the portion of the pipeline that contains input and output storage elements and that arbitrarily processes data stored in its storage elements.

Furthermore, the "device" ...valid data, but also when a stage requires more than one clock phase to finish processing its data. This also can occur when it creates valid data in one or both... ..control the passage of data between adjacent storage elements. The VALID signal may also be processed in an analogous manner.

A great advantage of the two-wire interface (one wire for... ..In addition, two extra latches and a small number of gates are preferably added to process the ACCEPT and VALID signals that are associated with the data latches in each half... ..application so requires. The interface in accordance with this embodiment can also be used to process analog signals.

As discussed previously, while other conventional timing arrangements may be used, the interface... ..circuit B1, which may be provided to convert output data from input latch LDIN into intermediate data, which is then later loaded in an output data latch LDOUT, which comprises the... ..is connected either directly as an input to the validation output latch LVOUT, or via intermediate logic devices or circuits that may alter the signal.

Similarly, the output validation signal QVOUT... ..to the input of the validation input latch QVIN of the following stage, or via intermediate devices or logic circuits, which may alter the validation signal. This ...word.

Preferred Data Structure - "tokens"

In the sample application shown in Fig. 4, each stage processes all input data, since there is no control circuitry that excludes any stage from allowing... ..are connected together in a relatively simple configuration. The simplest configuration is a pipeline of processing steps. For example, in the one shown in Fig. 1. The use of tokens, however...flows from left to right in the diagram. Data enters the machine and passes into processing Stage A. This may or may not modify the data and it then passes the... ..advantage of the tokens is their ability to achieve this kind of communication. Since any processing stage that does not recognize a token simply passes it on unaltered to the next... ..is transmitted along with the address and data fields in each token so that a processing stage can pass on a token (which can be of arbitrary length) without having to... ..be the first word of a new token.

Note that although the simple pipeline of processing stages is particularly useful, it will be appreciated that tokens may be applied to more complicated configurations of processing elements. An example of a more complicated processing element is described below.

It is not necessary, in accordance with the present invention, to... ..has extension bits. An example of this is a token that activates a stage that processes video quantization values stored in a quantization table (typically a memory device). For example, a...turn, is of great importance in video data pipeline systems since it ensures that all processing stages can be continuously running at full bandwidth.

In accordance with the present invention, in... ..some other chips in the set. This is advantageous both from the perspective of a customer and from that of a chip manufacturer. Even if modifications mean that all chips are...the end of a token (and hence the start of the next token) to be processed correctly (including simple non-manipulative transfer), even if the token is not recognized by the... ..block diagram of a pipeline stage whose function is as follows. If the stage is processing a predetermined token (known in this example as the DATA token), then it will duplicate... ..the address field of the DATA token. If, on the other hand, the stage is processing any other kind of token, it will delete every word. The overall effect is that... ..respective output signals:

In the duplication stage, the output from the data latch LDIN forms intermediate data referred to as MID(underscore)DATA. This intermediate data word is loaded into the data output latch LDOUT only when an intermediate acceptance signal (labeled "MID(underscore)ACCEPT" in Fig. 8a) is set HIGH.

The portion of... ..data. These include a "DATA(underscore)TOKEN" signal that indicates that the circuitry is currently processing a valid DATA Token, and a NOT(underscore)DUPLICATE signal which is used to control duplication of data. When the circuitry is processing a DATA Token, the NOT(underscore)DUPLICATE signal toggles between a HIGH and a LOW... ..the token to be duplicated once (but no more times). When the circuitry is not processing a valid DATA Token then the NOT(underscore)DUPLICATE signal is held in a HIGH state. Accordingly, this means that the token words that are being processed are not duplicated.

As Fig. 8a illustrates, the upper six bits of 8-bit intermediate data word and the output signal QI1 from the latch LI1 form inputs to a... ..explained further below.

Latch LO1 performs the function of latching the last value of the intermediate extension bit (labeled "MID(underscore)EXTN" and as signal S4), and it loads this value... ..and the DATA(underscore)TOKEN signal will become "0", indicating that the circuitry is not processing a DATA token.

If QI1 is "0" and SO is "0", thereby indicating a DATA... ..phase and the DATA(underscore)TOKEN signal will become "1", indicating that the circuitry is processing a DATA token.

The NOT(underscore)DUPLICATE signal (the output signal QO3) is similarly loaded...LVOUT at the same time that MID(underscore)DATA is loaded into LDOUT and the intermediate extension bit (signal S4) is loaded into

LEOUT. Signal S5 is also combined with the...above. This has the effect that all tokens except the one that causes the duplication process will be deleted from the token stream, since a device connected to the output terminals (OUTDATA, OUTEXTN and OUTVALID) will not recognize these token words as valid data.

As before... ..and is duplicated.

Referring now more particularly to Figure 10, there is shown a reconfigurable process stage in accordance with one aspect of the present invention.

Input latches 34 receive an... ..the input latches 34 is passed as a first input over line 35 to a processing unit 36. A first output from the token decode subsystem 33 is passed over line 37 as a second input to the

processing unit 36. A second output from the token decode 33 is passed over line 40 to an action identification unit 39. The action identification unit 39 also receives input from registers 43 and 44 over line 46. The registers 43... ..is determined by the history of tokens previously received. The output from the action identification unit 39 is passed over line 38 as a third input to the processing unit 36. The output from the processing unit 36 is passed to output latches 41. The output from the output latches 41 is... ..decoder 56 is passed over line 63 as an input to an Index to Data Unit (ITOD) 64. The Huffman decoder 56 and the ITOD 64 work together as a single logical unit. The output from the ITOD 64 is passed over line 65 to an arithmetic logic unit (ALU) 66. A first output from the ALU 66 is passed over line 67 to...blocks 133.

Referring to Figure 14b, in the JPEG and H.261 standards, the Common Intermediate Format (CIF) is used, wherein a picture 141 is encoded as 6 rows each containing... ..in a zigzag direction indicated by the arrow 144. The GOBs 142 are, in turn, processed row-by-row, left-to-right in each row.

Referring now to Figure 14c, it... ..in accordance with the practice of the present invention. A first picture 161 to be processed contains a first PICTURE(underscore)START token 162, first-picture information of indeterminate length 163, and a first PICTURE(underscore)END token 164. A second picture 165 to be processed contains a second PICTURE(underscore)START token 166, second picture information of indeterminate length 167... ..tokens 162 and 166 indicate the start of the pictures 161 and 165 to the processor. Likewise, the PICTURE(underscore)END tokens 164 and 168 signify the end of the pictures 161 and 165 to the processor. This allows the processor to process picture information 163 and 167 of variable lengths.

Referring to Figure 17, a split 171...Video Formatter (not shown in Figure 17).

Referring now to Figure 18, the prediction filtering process is illustrated. A forward picture 201 is passed over line 202 as a first input... ..the right of the value decode shift register 230, as indicated by area 231. This process eliminates overlapping start code images, as discussed below. A first output from the value decode... ..Code Detector. The Start Code Detector then receives a first data value image 244. Before processing the first data value image 244, the Start Code Detector may detect a second start... ..image 244 at a length 246. If this occurs, the Start Code Detector does not process the first data value image 244, and instead receives and processes a second data value image 247.

...line 1 of Table 600, whenever a "sequence start" image is received during H.261 processing or a "picture start" image is received during MPEG processing, the entire group of four control tokens is generated, each followed by its corresponding data...Picture Decoding

3. Motion Picture Decompression

4. RAM Memory Map

5. Bitstream Characteristics

6. Reconfigurable Processing Stage
7. Multi-Standard Coding
8. Multi-Standard Processing Circuit-2nd Mode of Operation
9. Start Code Detector
10. Tokens
11. DRAM Interface

12... ..described herein in greater detail) and reformatting this output for use, including display in a computer or other display systems, including a video display system. Implementation of this formatting varies significantly... ..the Spatial Decoder circuits.

The Spatial Decoder of the present invention performs all the required processing within a single picture. This reduces the redundancy within one picture.

The Temporal Decoder reduces... ..modeller 75, the inverse zig-zag 81 and the inverse DCT 83. The standard independent units within the Huffman decoder and parser include the ALU 66 and the token formatter 71.

Referring now to Figure 12, the standard-independent units include the DRAM interface 100, the fork 91, the FIFO register 96, the summer 98 and the output selector 106. The standard dependent units are the address generator 94, which is different in H.261 and in MPEG, and...much of the operation is very similar between the three different compression standards.

The next unit is the state machine 68 (Figure 11) located within the Huffman decoder and parser. Here... ..The same holds true for JPEG, which is a third,completely independent program.

The next unit is the Huffman decoder 56 which functions with the index to data unit 64. Those two units cooperate together to perform the Huffman decoding. Here, the algorithm that is used for Huffman... ..to the Huffman decoder at different times consistent with the standard in operation.

The last unit on the chip that is dependent on the compression standard is the inverse quantizer 79...an H.261 group of blocks and an MPEG slice. When H.261 data is processed after the Start Code Detector, each group of blocks is preceded by a slice(underscore... ..these standards have totally different sets of tables.

As previously indicated, most of the system units are compression standard independent. If a unit is standard independent, and such units need not remember what CODING(underscore)STANDARD is being processed. All of the units that are standard dependent remember the compression standard as the CODING(underscore)STANDARD token flows... ..CODING(underscore)STANDARD tokens at the Start Code Detector that is positioned as the first unit in the pipeline, this change of compression standard is readily handled. The token says a... ..found in the standard, i.e. from the bitstream into a prediction mode token. This processing is performed by the Huffman decoder and parser state machine, where it is easy to... ..to that token. By having these tokens and using them appropriately, the design of other units in the machine is simplified. Although there may be some complications in the program, benefits... ..a first encoded signal (the MPEG or H.261 encoded video signal) in a pipeline processing system. The Temporal Decoder is not needed for JPEG decoding.

In this regard, the invention... ..the use of a single pipeline decoder and decompression system. The decoding and decompression pipeline processor is organized on a unique and special configuration which allows the handling of the multi... ..video signals through the use of techniques all compatible with the single pipeline decoder and processing system. The Spatial Decoder is combined with the Temporal Decoder, and the Video Formatter is...with only still pictures. The compression standard independent Spatial Decoder performs all of the data processing within the boundaries of a single picture. Such a decoder handles the spatial decompression of... ..to the multi-standard,

configurable Video Formatter, which then provides an output to the display terminal. In a first sequence of similar pictures, each decompressed picture at the output of the... of control tokens and DATA tokens, in combination with a plurality of sequentially-positioned reconfigurable processing stages selected and organized to act as a standard-independent, reconfigurable-pipeline-processor.

With regard to JPEG decoding, a single Spatial Decoder with no off chip DRAM can... video. Accordingly, signals carried by DATA tokens pass directly through the Temporal Decoder without further processing when the Temporal Decoder is configured for a JPEG operation.

Another aspect of the present... for subsequent use in temporal decoding of subsequent pictures.

Generally, the Temporal Decoder performs the processing between pictures either earlier and/or later in time with reference to the picture currently... is distributed among several areas of DRAM in the sense that the decompressed output information, processed by the Spatial Decoder, is stored in other DRAM registers by other random access memories... first decoder circuit (the Spatial Decoder) directly to the Video Formatter for handling without signal processing delay.

The Temporal Decoder also reorders the blocks of picture data for display by a... from a selection of pictures which have arrived earlier or later than the picture under processing. When a picture is described in this context, it may mean any one of the... 2. The result, i.e., the final decoded picture resulting from the addition of a process step performed by the decoder;

3. Previously decoded pictures read from the DRAM; and

4... START token and a subsequent PICTURE(underscore)END token.

After the picture data information is processed by the Temporal Decoder, it is either displayed or written back into a picture memory location. This information is then kept for further reference to be used in processing another different coded data picture.

Re-ordering of the MPEG encoded pictures for visual display... used to encode a referenced picture of a picture might be identified as being one unit long, another picture might be a number of units long, while still a third picture could be a fraction of that unit.

None of the existing standards (MPEG 1.2, JPEG, H.261) define a way of... picture rate, whereas the Video Formatter can handle a variable input picture rate.

6. RECONFIGURABLE PROCESSING STAGE

Referring again to Figure 10, the reconfigurable processing stage (RPS) comprises a token decode circuit 33 which is employed to receive the tokens... input latches 34. The output of the token decode circuit 33 is applied to a processing unit 36 over the two-wire interface 37 and an action identification circuit 39. The processing unit 36 is suitable for processing data under the control of the action identification circuit 39. After the processing is completed, the processing unit 36 connects such completed signals to the output, two-wire interface bus 40 through output... token decode circuit 33 are applied simultaneously to the action identification circuit 39 and the processing unit 36. The action identification function as well as the RPS is described in further detail... not standard independent circuits. The data flows through the token decode circuit 33, through the

processing unit 36 and onto the two-wire interface circuit 42 through the output latches 41. If... wire interface 42 through the output circuit 41. The present invention operates as a pipeline processor having a two-wire interface for controlling the movement of control tokens through the pipeline... time, the token decode circuit 33 provides a proper flag or index signal to the processing unit 36 to alert it to the presence of the token being handled by the action identification circuit 39.

Control tokens may also be processed.

A more detailed description of the various types of tokens usable in the present invention...standard now passing through the state machine shown with reference to Figure 10.

Similarly, the processing unit 36 which is under the control of the action identification circuit 39 is now ready to process the information contained in the data fields of the DATA token when it is appropriate... ..action identification circuit 39 and is immediately followed by a DATA token which is then processed by the processing unit 36. The control token exits the output latches circuit 41 over the output two-wire interface 42 immediately preceding the DATA token which has been processed within the processing unit 36.

In the present invention, the action identification circuit, 39, is a state machine holding... ..show that the action can also be affected by the token that is currently being processed by the token decode circuit 33.

In general, there is shown token decoding and data processing in accordance with the present invention. The data processing is performed as configured by the action identification circuit 39. The action is affected by... ..information stored from previously decoded tokens in registers 43 and 44, the current token under processing, and the state and history information that the action identification unit 39 has itself acquired. A distinction is thereby shown between Control tokens and DATA tokens.

In any RPS, some tokens are viewed by that RPS unit as being Control tokens in that they affect the operation of the RPS presumably at... ..are viewed by the RPS as DATA tokens. Such DATA tokens contain information which is processed by the RPS in a way that is determined by the design of the particular... ..view of the same token. Some of the tokens might be viewed by one RPS unit as DATA Tokens while another RPS unit might decide that it is actually a Control Token. For example, the quantization table information... ..into a token called a quantization table token (QUANT(underscore)TABLE) which goes down the processing pipeline. As far as that machine is concerned, all of that was data; it was... ..sort of data into another sort of data, which is clearly a function of the processing performed by that portion of the machine. However, when that information gets to the inverse... ..present. This information is viewed as control information, and then that control information affects the processing that is done on subsequent DATA tokens because it affects the number that you multiply... ..important feature of the invention is that each of the stages of circuitry has the processing capability within it to be able to perform the necessary operations for each of the... ..operations are to be performed at a given time, come as tokens. There is one processing element that differs between the different stages to provide this capability. In the state machine...standard is and it looks up the parameters that it needs to apply to the processing elements in order to perform a proper operation. For example, the inverse quantizer will look... ..is set to 1 for a particular compression standard, and will apply that to its processing circuitry.

In a similar sense the Huffman decoder 56 has a number of tables within... ..MPEG video standard or the JPEG video standard. These three compression coding standards specify similar processes to be done on the arriving data, but the structure of the datastreams is different... ..token stream embodying the current coding standard. The control tokens are passed through the pipeline processor, and are used, i.e., decoded, in the state machines to which they are relevant... ..this regard, the DATA Tokens are treated in the same fashion, insofar as they are processed only in the state machines that are configurable by the control tokens into processing such DATA Tokens. In the remaining state machines, they pass through unchanged.

More specifically, a... ..signals. The remaining portions of the token are used to indicate and identify the internal processing control function which is standard for all of the datastreams passing through the pipeline processor. In one form of the invention, the token extension is used to carry the current...accompanying data. As previously discussed, this information is utilized in the system to reconfigure the processing stage used to perform the function required by the various standards created for that purpose... ..picture number as indicated by the value.

The system also includes a multi-stage parallel processing pipeline operating under the principles of the two-wire interface previously described. Each of the... the token presently entering the state machine into the action identification circuit 39 or the processing unit 36, as appropriate. The processing unit has been previously reconfigured by the next previous control token into the form needed for handling the current coding standard, which is now entering the processing stage and carried by the next DATA token. Further, in accordance with this aspect of the invention, the succeeding state machines in the processing pipeline can be functioning under one coding standard, i.e., H.261, while a previous... tokens required to decode a number of coding standards with a fixed number of reconfigurable processing stages. More specifically, the PICTURE(underscore)END control token is employed because it is important... standard machine, it is necessary to create additional control tokens within the multi-standard pipeline processing machine which will then indicate which one of the standard decoding techniques to use. Such... and to push the current picture through the decoder to the display.

8. MULTI-STANDARD PROCESSING CIRCUIT - SECOND MODE OF OPERATION

A compression standard-dependent circuit, in the form of the... of the Start Code Detector will subsequently be discussed in further detail, as will the process of starting up of the decoder.

The aforementioned description has been concerned primarily the data which immediately follows according to the standard. However, in the multi-standard pipeline processing system of the present invention, where compatibility is required for multiple standards, the system has... signals, including flag signals, are generated by each state machine to handle some of the processing within that state machine. Values carried in the standards can be used to access machine... its contents must be removed from the two wire interface to ensure that no further processing takes place using these 3 bytes. The decode register is emptied, and the value decode... 10. TOKENS

In the practice of the present invention, a token is a universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions and is adapted for use with a reconfigurable processing stage (RPS) which is a stage, which in response to a recognized token, reconfigures itself to perform various operations.

Tokens may be either position dependent or position independent upon the processing stages for performance of various functions. Tokens may also be metamorphic in that they can be altered by a processing stage and then passed down the pipeline for performance of further functions. Tokens may interact... other functions, and the specific interaction with a stage may be conditioned by the previous processing history of a stage.

A PICTURE(underscore)END token is a way of signalling the... through a fixed size, fixed width buffer.

The present invention is directed to a pipeline processing system which has a variable configuration which uses tokens and a two-wire system. The... do not use control tokens.

The control tokens are generated by circuitry within the decoder processor and emulate the operation of a number of different type standard-dependent signals passing into the serial pipeline processor for handling. The technique used is to study all the parameters of the multi-standards that are selected for processing by the serial processor and noting 1) their similarities, 2) their dissimilarities, 3) their needs and requirements and 4) selecting the correct token function to effectively process all of the standard signals sent into the serial processor. The functions of the tokens are to emulate the standards. A control token function is... the standard dependent signals and as an element to transmit control information through the pipeline processor.

In prior art system, a dedicated machine is designed according to well-known techniques to... tokens provide and make a sensible format for communicating information through the decompression circuit pipeline processor. In the design selected hereinafter and used in the preferred embodiment, each word of a... However, this is not a limitation on the invention, but on the magnitude of the processing steps elected to be accomplished by use of these tokens. It is to be noted... bit address for use in accessing the random access memories used throughout this serial decompression processor. This provides an additional degree of variability that facilitates a broad range of versatility.

As previously described, the DATA token carries data from one processing stage to the next. Consequently, the characteristics of this token change as it passes through... ..longest number of data bits because it needs to provide the most information to the

processing unit so that it can start the decompression with as much information as possible. Words which...to receive an address, it waits for the address generator to supply a valid address, processes that address and then sets the accept line high for one clock period. Thus, it... ..be read. This signal passes between two asynchronous clock regimes and, therefore, passes through three synchronizing flip flops.

Provided RAM2 312 is empty, the next item of data to arrive on...interesting.

In general, prediction data will be offset from the position of the block being processed as specified in the motion vectors in x and y. Thus, the block of data... ..address, 9. Data is read from this address and the x value is incremented. The process is repeated until the x value reaches its stop value, at which point, the y... ..is read, the x value is again incremented until it reaches its stop value. The process is repeated until both x and y values have reached their stop values. Thus, the...invention, is that additional information must be provided to the prediction filters to indicate what processing is required on the data. This consists of the following:

a "last byte" signal indicating... ..bit 0) is incremented and the x address (3 LSBS) is reset to zero. This process is repeated until 64 bytes have been read. With a 16 or 32 bit wide...register while its access register is set to zero, the results are undefined.

14. MICRO-PROCESSOR INTERFACE

A standard byte wide micro-processor interface (MPI) is used on all circuits with in the Spatial Decoder and Temporal Decoder... ..the parameter column. The actual specifications are shown in the respective columns min, max and units.

The DC operating conditions can be seen with reference to Table A.6.3. Here... ..the signal is present the maximum amount of time that this signal is available. The Units column gives the units of measurement used to describe the signals.

16. MPI WRITE TIMING

The general description of...a PICTURE(underscore)END token is decoded and forces the data in the coded data buffers to be applied to the Huffman decoder and video demultiplexor, the final picture can be... ..Consequently, the machine will not go into error recovery mode and will successfully continue to process the coded data.

A still further advantage of the use of a PICTURE(underscore)END token is that the serial pipeline processor will continue the processing of uninterrupted data. Through the use of a PICTURE(underscore)END token, the serial pipeline processor is configured to handle less than the expected amount of data and, therefore, continues processing. Typically, a prior art machine would stop itself because of an error condition. As previously... ..of the Huffman decode and Video Demultiplexor know the number of blocks that it will process during each picture recovery cycle. When the correct number of blocks do not arrive from... ..Each of the state machines recognizes a FLUSH control token as information not to be processed. Accordingly, the FLUSH token is used to fill up all of the remaining empty parts... ..Huffman Decoder and Video Demultiplexor. In this way, the FLUSH token is like padding for buffers.

The Token Decoder in the Huffman circuit recognizes the FLUSH token and ignores the pseudo... ..less information than normally expected to decode the last picture. The Huffman decode circuit finishes processing the information contained in the last picture, and outputs this information through the DRAM interface... ..token, in accordance with the present invention, is used to pass through the entire pipeline processor and to ensure that the buffers are emptied

and that other circuits are reconfigured to... ..END token, a padding word and a FLUSH token indicating to the serial pipeline processor that the picture processing for the current picture form is completed. Thereafter, the various state machines need reconfiguring to... ..FLUSH token resets each stage as it passes through, but-allows subsequent stages to continue processing. This prevents a loss of data. In other words, the FLUSH token is a variable... ..AFTER PICTURE

The STOP(underscore)AFTER(underscore)PICTURE function is employed to shut down the processing of the serial pipeline decompressing circuit at a logical point in its operation. At this... ..a picture, the STOP(underscore)AFTER(underscore)PICTURE operation signals the end of all current processing.

22. MULTI(underscore)STANDARD - SEARCH MODE

Another feature of the present invention is the use... ..MODE control token which is used to reconfigure the input to the serial pipeline processor to look at the incoming bit stream. When the search mode is set, the Start... ..combination of control tokens, and DATA tokens along with the reconfiguration circuits, to provide similar processing.

The use of search mode in the present invention is convenient in many situations including... ..video disc. In general, a search mode is convenient when the user interrupts the normal processing of the serial pipeline at a point where the machine does not expect such an...be the case.

In brief, the Huffman Decoder 321 works in conjunction with the other units shown in Figure 27. These other units are the Parser State Machine 322, the inshifter 323, the Index to Data unit 324, the ALU 325, and the Token Formatter 326. As described previously, connection between these blocks is governed by a two wire interface. A more detailed description of how these units function is subsequently described herein in greater detail, the focus here is on particular aspects... ..control certain functions of the Index to Data 324 and ALU 325. Control of these units by the Huffman Decoder is necessary for proper decoding of block-level information. Having the... ..further detail in the "More Detailed Description of the Invention" section.

The Index to Data unit 324 performs the second part of the multi-part algorithm. This unit contains a look up table that provides the actual Huffman decoded data. Entries in the...by detecting these in the Huffman Decoder 321, rather than in the Index to Data unit 324.

This index number is then passed to the Index to Data unit 324. In essence, the Index to Data unit is a look-up table. In accordance with one aspect of the algorithm, the look... ..format that JPEG specifies for transferring an alternate JPEG table.

From the Index to Data unit 324, the decoded index number or other data is passed, together with the accompanying control... ..the entering data to ensure that the DATA tokens are of the correct size for processing. In fact, the token stream can be corrected in some situations if the error is... ..an order that is useful for the decompression circuits, but not for the particular display unit being used. When a block of data enters the Buffer Manager, the Buffer Manager supplies...the output of the Spatial Decoder or Temporal Decoder and re-format it for a computer or display system. The details of this formatting will vary between applications. In a simple...Token. The DATA Token can have as many bits as are necessary for carrying out processing at a particular place in the system. All other Tokens ignore the extra bits.

A.3.2 The DATA Token

The DATA Token carries data from one processing stage to the next. Consequently, the characteristics of this Token change as it passes through... ..will be sufficient to collect DATA Tokens and to detect a few Tokens that provide synchronization information (such as PICTURE(underscore)START). In this regard, see subsequent sections A.16, "Connecting... ..from the data stream. This provides an alternative to doing the configuration via the micro processor interface.

A.3.4 Description of Tokens

This section documents the Tokens which are implemented... 3.5.1. Note: JPEG requires a 2:1:1 structure for its macroblocks when processing 4:2:2 data. See Table A.3.5.

A.3.6 Special Token formats...either is low then the interface is taken to high impedance.

Note: on-chip data processing is not terminated when the DRAM interface is at high impedance. Therefore, errors will occur...decoded video's picture rate. Accordingly, this clock can be used to provide audio/video synchronization.

A.7.1 Spatial Decoder clock signals

The Spatial Decoder has two different (and potentially... in accordance with the present invention, must know what video standard is being input for processing. Thereafter, the system can accept either pre-existing Tokens or raw byte data which is...time a value is written into coded(underscore)data (7:0). Software is responsible for settling coded(underscore)extn to 0 before the last word of any Token is written to... 0). The start of this new DATA Token then passes into the Spatial Decoder for processing.

Each time a new 8 bit value is written to coded(underscore)data (7:0... Detector analyses data in the DATA Tokens bit serially. The Detector's normal rate of processing is one bit per clock cycle (of coded(underscore)clock). Accordingly, it will typically decode a byte of coded data every 8 cycles of coded(underscore)clock. However, extra processing cycles are occasionally required, e.g., when a non-DATA Token is supplied or when the main decoder(underscore)clock. Data transfer is synchronized to decoder(underscore)clock on-chip.

SECTION A.11 Start code detector

A.11.1... Code Detector. So, accessing these registers will be unreliable if the Start Code Detector is processing data. The user is responsible for ensuring that the Start Code Detector is halted before... Detector. In this case, the Tokens are passed through the Start Code Detector with no processing to other stages of the Spatial Decoder. These Tokens can only be inserted just before...result will be unpredictable if this is done when the Start Code Detector is actively processing data.

Discard all mode can be safely initiated after any of the Start Code Detector...start code non-alignment interrupt is suppressed.

In contrast, however, JPEG was designed for a computer environment where byte alignment is guaranteed. Therefore, marker codes should only be detected when byte... the other hand, was designed to meet the needs of both communications (bit serial) and computer (byte oriented) systems. Start codes in MPEG data should normally be byte aligned. However, the...result will be unpredictable if this is done when the Start Code Detector is actively processing data. So, before initiating a start code search, the Start Code Detector should be stopped so no data is being processed. The Start Code Detector is always in this condition if any of the Start Code...the spatial video decoding circuits (inverse modeler, quantizer and DCT). This second logical buffer allows processing time to include a spread so as to accommodate processing pictures having varying amounts of data.

Both buffers are physically held in a single off 1.1, the unit for all the above mentioned registers is a 512 bit block of data. Accordingly, the... until there is space in the buffer. If a buffer continues to be full, more processing stages "up steam" of the buffer will halt until the Spatial Decoder is unable to... converting coded data into Tokens started by the Start Code Detector. There are four main processing blocks in the Video Demux: Parser State Machine, Huffman decoder (including an ITOD), Macroblock counter... or state machine follows the syntax of the coded video data and instructs the other units. The Huffman decoder converts variable length coded (VLC) data into integers. The Macroblock counter keeps...

Claims:

1. An apparatus for providing a time delay to a group of compressed pictures, the pictures corresponding to... ..data from said buffer and capable of delaying said words of data ;
a control circuit intermediate and in communication with said counter circuit and said inverse modeller circuit ;
said counter circuit... ..have met a start-up criterion and controlling said inverse modeller delay feature.
2. The apparatus as recited in claim 1, having a data stream including run level code, further characterized... ..a level, whereby each token is expressed with a specified number of values.
3. The apparatus as recited in any one of claims 1 and 2, wherein said token is a DATA token.
4. The apparatus as recited in any one of claims 1 to 3, wherein said inverse modeller means blocks tokens which lack said specified number of values.
5. The apparatus as recited in any one of claims 1 to 4, wherein said specified number of values is 64 coefficients.
6. The apparatus as recited in claim 1, wherein said token is a QUANT(underscore)TABLE token for causing said processing stage to generate a quantization table.

6/K/9 (Item 9 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Start code detecting apparatus for video data stream

Country	Number	Kind	Date
---------	--------	------	------

Abstract ...a video decoding and decompression system having an input, an output and a plurality of processing stages between the input and the output defining a pipeline, the improvement comprising :

Type	Pub. Date	Kind	Text
Available Text	Language	Update	Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: INTRODUCTION

The present invention is directed to improvements in methods and apparatus for decompression which operates to decompress and/or decode a plurality of differently encoded input... ..of the well known standards known as JPEG, MPEG and H.261.

A serial pipeline processing system of the present invention comprises a single two-wire bus used for carrying unique... ..to a plurality of adaptive decompression circuits and the like positioned as a reconfigurable pipeline processor.

PRIOR ART

One prior art system is described in United States Patent No. 5,216,724. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and a host processor. The device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus.

United States Patent No. 4,785... ...a known quad tree data structure.

United States Patent No. 5,122,875 discloses an apparatus for encoding/decoding an HDTV signal. The apparatus includes a compression circuit responsive to high definition video source signals for providing hierarchically layered... ...compressed video data of relatively greater and lesser importance to image reproduction respectively. A transport processor, responsive to the high and low priority codeword sequences, forms high and low priority transport... ...United States Patent No. 5,168,356 discloses a video signal encoding system that includes apparatus for segmenting encoded video data into transport blocks for signal transmission. ...in respective transport blocks.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening. This is accomplished by an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ...frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ...oscillators and the receiver can continuously receive each channel, then the receiver need not be synchronized with the transmitter. An FFT algorithm implements a fast discrete approximation to the continuous case in which the receiver synchronizes to the first frame and then acquires subsequent frames every frame period. The frame period... ...increasing the amount of data transmitted.

United States Patent No. 5,212,742 discloses an apparatus and method for processing video data for compression/decompression in real-time. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and host processor. Lastly, the device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus. The method handles assigning portions of the image for each of the processors to operate upon.

United States Patent No. 5,231,484 discloses a system and method... ...MPEG standards. Included are three cooperating components or subsystems that operate to variously adaptively pre-process the incoming digital motion video sequences, allocate bits to the pictures in a sequence, and... ...States Patent No. 5,267,334 discloses a method of removing frame redundancy in a computer system for a sequence of moving images. The method comprises detecting a first scene change... ...facing" keyframe or intraframe, and it is normally present in CCITT compressed video data. The process then comprises generating at least one intermediate compressed frame, the at least one intermediate compressed frame containing difference information from the first image for at least one image following...change, known as a "backward-facing" keyframe. The first keyframe and the at least one intermediate compressed frame are linked for forward play, and the second keyframe and the intermediate compressed frames are linked in reverse for reverse play. The intraframe may also be used... ...of complete scene information.

United States Patent No. 5,276,513 discloses a first circuit apparatus, comprising a given number of prior-art image-pyramid stages, together with a second circuit apparatus, comprising the same given number of novel motion-vector stages, perform cost-effective hierarchical motion analysis (HMA) in real-time, with minimum system processing delay and/or employing minimum system processing delay and/or employing minimum hardware structure. Specifically, the first and second circuit apparatus, in response to relatively high-resolution image data from an ongoing input series of successive... ...a relatively high frame rate (e.g., 30 frames per second), derives, after

a certain processing-system delay, an ongoing output series of successive given pixel-density vector-data frames that... ..of successive image frames.

United States Patent No. 5,283,646 discloses a method and apparatus for enabling a real-time video encoding system to accurately deliver the desired number of... ..desired bit allocations.

The article, Chong, Yong M., A Data-Flow Architecture for Digital Image Processing, Wescon Technical Papers: No. 2 Oct./Nov. 1984, discloses a real-time signal processing system specifically designed for image processing. More particularly, a token based data-flow architecture is disclosed wherein the tokens are of... ..width having a fixed width address field. The system contains a plurality of identical flow processors connected in a ring fashion. The tokens contain a data field, a control field and a tag. The tag field of the token is further broken down into a processor address field and an identifier field. The processor address field is used to direct the tokens to the correct data-flow processor, and the identifier field is used to label the data such that the data-flow processor knows what to do with the data. In this way, the identifier field acts as an instruction for the data-flow processor. The system directs each token to a specific data-flow processor using a module number (MN). If the MN matches the MN of the particular stage... ..to locate the decoder in the preceding stage in order to pre-decode complex decoding processing and to alleviate critical path problems in the logic circuit. The elastic nature of the...of block signal in most cases.

United States Patent No. 4,903,018 discloses a process and data processing system for compressing and expanding structurally associated multiple data sequences. The process is particular to data sets in which an analysis is made of the structure in... ..data series on the basis of the order number of these data elements. The data processing system for performing the processes includes a storage matrix (26) and an index storage (28) having line addresses of the... ..the final actual video.

United States Patent No. 5,060,242 discloses an image signal processing system DPCM encodes the signal, then Huffman and run length encodes the signal to produce... ..tightly packed without gaps for efficient transmission without loss of any data. The tightly packed apparatus has a barrel shifter with its shift modulus controlled by an accumulator receiving code word... ..OR gate is connected to the shifter, while a register is connected to the gate. Apparatus for processing a tightly packed and decorrelated digital signal has a barrel shifter and accumulator for unpacking... ..an inverse DCPM decoder.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening is accomplished by use of an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ..frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ..kernel matrix.

United States Patent No. 5,231,486 discloses a high definition video system processes a bitstream including high and low priority variable length coded Data words. The coded Data... ..packed High Priority Data and packed Low Priority Data by means of respective data packing units. The coded Data is continuously applied to both packing units. High Priority and Low Priority Length words indicating the bit lengths of high priority and... ..States Patent No. 5,287,178 discloses a video signal encoding system includes a signal processor for segmenting encoded video data into transport blocks having a header section and a packed data section. The system also includes reset control apparatus for releasing resets of system components, after a global system reset, in a prescribed non-simultaneous phased sequence to enable signal processing to commence in the prescribed sequence. The phased reset release sequence begins when valid data...United States Patent No. 5,142,380 to Sakagami et al. discloses an image compression apparatus

suitable for use with still images such as those formed by electronic still cameras using... ..and Q.

United States Patent No. 5,193,002 to Guichard et al. disclosed an apparatus for coding/decoding image signals in real time in conjunction with the CCITT standard H.261. A digital signal processor carries out direct quantization and reverse quantization.

United States Patent No. 5,241,383 to Chen et al. describes an apparatus with a pseudo-constant bit rate video coding achieved by an adjustable quantization parameter. The... ..relates to an improved pipeline system having an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along the pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the processing stages in the pipeline and interacting with selected stages in the pipeline for control data and/or combined control-data functions among the processing stages, so that the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing. In accordance with the invention, the processing stages may be configurable in response to recognition of at least one token. One of the processing stages may be a Start Code Detector which receives the input and generates and/or... ..and resetting the system, and a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards. The present invention... ..data and having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video pictures of varying data size can be controlled. Also in accordance with the invention, a processing stage receives the input data stream, the stage including means for recognizing specified bit stream patterns, whereby the processing stage facilitates random access and error recovery. The invention may also include a means for... ..invention also includes an inverse modeller stage, an inverse discrete cosine transform stage, and a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for processing data.

In addition, the present invention relates to an improved pipeline system having a Huffman...pipeline stage that incorporates a two-wire transfer control and also shows two consecutive pipeline processing stages with the two-wire transfer control;

Figures. 5a and 5b taken together depict one... ..shown in Figures. 8a and 8b.

Figure 10 is a block diagram of a reconfigurable processing stage;

Figure 11 is a block diagram of a spatial decoder;

Figure 12 is a... ..decoder including the prediction filters;

Figure 18 is a pictorial representation of the prediction filtering process;

Figure 19 shows a generalized representation of the macroblock structure;

Figure 20 shows a generalized... ..buffer;

Figure 25 is a pictorial diagram illustrating prediction data offset from the block being processed;

Figure 26 is a pictorial diagram illustrating prediction data offset by (1,1);

Figure 27...in general terms, the present invention provides an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along a pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the stages in the pipeline and interacting with selected stages in the pipeline for control, data and/or combined control-data functions among the processing stages, whereby the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing.

Each of the processing stages in the pipeline may include both primary and secondary storage, and the stages in processing stages for performance of functions or position independent of the processing stages for performance of functions.

In a pipeline machine, in accordance with the invention, the... ..altered by interfacing with the stages, and the tokens may interact with all of the processing stages in the pipeline or only with some but less than all of said processing stages. The tokens in the pipeline may interact with adjacent processing stages or with non-adjacent processing stages, and the tokens may reconfigure the processing stages. Such tokens may be position dependent for some functions and position independent for other... ..be Huffman coded.

In the improved pipeline machine, the tokens may be generated by a processing stage. Such pipeline tokens may include data for transfer to the processing stages or the tokens may be devoid of data. Some of the tokens may be identified as DATA tokens and provide data to the processing stages in the pipeline, while other tokens are identified as control tokens and only condition the processing stages in the pipeline, such conditioning including reconfiguring of the processing stages. Still other tokens may provide both data and conditioning to the processing stages in the pipeline. Some of said tokens may identify coding standards to the processing stages in the pipeline, whereas other tokens may operate independent of any coding standard among the processing stages. The tokens may be capable of successive alteration by the processing stages in the pipeline.

In accordance with the invention, the interactive flexibility of the tokens in cooperation with the processing stages facilitates greater functional diversity of the processing stages for resident structure in the pipeline, and the flexibility of the tokens facilitates system... ..or alteration. The tokens may be capable of facilitating a plurality of functions within any processing stage in the pipeline. Such pipeline tokens may be either hardware based or software based... ..system bandwidth in the pipeline. The tokens may provide data and control simultaneously to the processing stages in the pipeline.

The invention may include a pipeline processing machine for handling plurality of separately encoded bit streams arranged as a single serial bit... ..and for passing unrecognized control tokens along the pipeline, and a reconfigurable decode and parser processing means responsive to a recognized control token for reconfiguring a particular stage to handle an... ..be a pipeline system and the Start Code Detector may be positioned as the first processing stage in the pipeline.

The present invention also provides, in a system having a plurality of processing stages, a universal adaptation unit in the form of an interactive interfacing token for control and/or data functions among the processing stages, the token being a PICTURE(underscore)START code token for indicating that the start... ..The token may also be a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards.

The CODING(underscore... ..standard as JPEG, and/or any other appropriate picture standard. At least some of the processing stages reconfigure in response to the CODING(underscore)STANDARD token.

One of the processing stages in the system may be a Huffman decoder and parser and, upon receipt of... ..Data stage, and the parser stage may send an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... ..video data, having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video... ..controlled.

The system may include a spatial decoder having a two-wire interface interconnecting processing stages, the interface enabling serial processing for data and parallel processing for control.

As previously indicated, the system may further include a ROM having separate stored... ..of a plurality of picture standards, the programs being selectable by a token to facilitate processing for a plurality of different picture standards.

The spatial decoder system also includes a token... ..decoding stage and a parser stage for sending an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... ..The present invention also provides a pipeline system having an input data stream, and a processing stage for receiving the input data stream, the stage including means for recognizing specified bit... ..whereby said stage facilitates random access and error recovery. In accordance with the invention, the processing stage may be a start code detector and the bit stream patterns may include start... ..token and padding insures uniformity of word size. In accordance with the invention, a reconfigurable processing stage may be provided as a spatial decoder and the padding means adds to picture... ..that if the DATA token has less than the predetermined length, the padder circuit adds units of data to the DATA token until the predetermined length is achieved. A bypass circuit...1 tokens into a buffer, having a second predetermined width.

The invention also provides an apparatus for providing a time delay to a group of compressed pictures, the pictures corresponding to... ..and capable of delaying the words of data, is in communication with a control circuit intermediate the counter circuit and the inverse modeller circuit, the control circuit also communicating with the... ..inverse modeller stage and an inverse discrete cosine transform stage, the improvement characterized by a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for

processing data.

In accordance with the invention, the token may be a QUANT(underscore)TABLE token for causing the processing stage to generate a quantization table.

The present invention also provides a Huffman decoder for... ..of bits used to represent an item of data.

DECODER: An embodiment of a decoding process.

DECODING (PROCESS): The process defined in this specification that reads an input coded bitstream and produces decoded pictures or... ..the same order in which they were presented at the input of the encoder.

ENCODING (PROCESS): A process, not specified in this specification, that reads a stream of input pictures or audio samples...to provide an estimate of the pel value or data element currently being decoded.

RECONFIGURABLE PROCESS STAGE (RPS): A stage, which in response to a recognized token, reconfigures itself to perform various operations.

SLICE: A series of macroblocks.

TOKEN: A universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions... ..indicates that the corresponding stage holds valid data, i.e., data that is to be processed in one of the pipeline stages. After processing (which may involve nothing more than a simple transfer without manipulation of the data) valid... ..present invention may be used with any number of pipeline stages. Furthermore, data may be processed in more than one stage and the processing time for different stages can differ.

In addition to clock and data signals (described below... ..other system. For example, the last pipeline stage may pass its data on to subsequent processing circuitry. The ACCEPT signal, which is illustrated as the lower of the two lines connecting... ..the minimum disturbance possible to other pipeline stages. Succeeding pipeline stages are allowed to continue processing and, therefore, this means that gaps open up in the stream of data following the... ..The data in the pipeline is encoded such that many different types of data are processed in the pipeline. This encoding accommodates data packets of variable size and the size of...the other hand, it may generate itself, all or part of the data to be processed in the pipeline. Indeed, as is explained below, a "stage" may contain arbitrary

processing circuitry, including none at all (for simple passing of data) or entire systems (for example... values zero and 255 may not be used.

If such a picture were to be processed in a pipeline built in the practice of the present invention, then one of these... data must not be written over since it is data that must be saved for processing or use in a downstream device e.g., a pipeline stage, a device or a... connected to the pipeline upstream contains data D4 that is to be transferred into and processed in the pipeline. ...pipeline, in accordance with the preferred embodiments of the present invention, to "fill up" empty processing stages is highly advantageous since the processing stages in the pipeline thereby become decouple from one another. In other words, even though... data can be transferred into the pipeline and between stages even when one or more processing stages is blocked.

In the embodiment shown in Fig. 1, it is assumed that the... propagate all the way back to the beginning of the pipeline if there is some intermediate stage that is able to accept new data.

In the embodiment illustrated in Fig. 1... has been mentioned. It is to be further understood that each pipeline stage may also process the data it has received arbitrarily before passing it between its internal storage elements or... the portion of the pipeline that contains input and output storage elements and that arbitrarily processes data stored in its storage elements.

Furthermore, the "device" ... valid data, but also when a stage requires more than one clock phase to finish processing its data. This also can occur when it creates valid data in one or both... control the passage of data between adjacent storage elements. The VALID signal may also be processed in an analogous manner.

A great advantage of the two-wire interface (one wire for... In addition, two extra latches and a small number of gates are preferably added to process the ACCEPT and VALID signals that are associated with the data latches in each half... application so requires. The interface in accordance with this embodiment can also be used to process analog signals.

As discussed previously, while other conventional timing arrangements may be used, the interface... circuit B1, which may be provided to convert output data from input latch LDIN into intermediate data, which is then later loaded in an output data latch LDOUT, which comprises the... is connected either directly as an input to the validation output latch LVOUT, or via intermediate logic devices or circuits that may alter the signal.

Similarly, the output validation signal QVOUT... to the input of the validation input latch QVIN of the following stage, or via intermediate devices or logic circuits, which may alter the validation signal. This ...word.

Preferred Data Structure - "tokens"

In the sample application shown in Fig. 4, each stage processes all input data, since there is no control circuitry that excludes any stage from allowing... are connected together in a relatively simple configuration. The simplest configuration is a pipeline of processing steps. For example, in the one shown in Fig. 1. The use of tokens, however... flows from left to right in the diagram. Data enters the machine and passes into processing Stage A. This may or may not modify the data and it then passes the... advantage of the tokens is their ability to achieve this kind of communication. Since any processing stage that does not recognize a token simply passes it on unaltered to the next... is transmitted along with the address and data fields in each token so that a processing stage can pass on a token (which can be of arbitrary length) without having to... be the first word of a new token.

Note that although the simple pipeline of processing stages is particularly useful, it will be appreciated that tokens may be applied to more complicated configurations of processing elements. An example of a more complicated processing element is described below.

It is not necessary, in accordance with the present invention, to... has extension bits. An example of this is a token that activates a stage that processes video quantization values stored in a quantization table (typically a memory

device). For example, a...turn, is of great importance in video data pipeline systems since it ensures that all processing stages can be continuously running at full bandwidth.

In accordance to the present invention, in... some other chips in the set. This is advantageous both from the perspective of a customer and from that of a chip manufacturer. Even if modifications mean that all chips are...the end of a token (and hence the start of the next token) to be processed correctly (including simple non-manipulative transfer), even if the token is not recognized by the... block diagram of a pipeline stage whose function is as follows. If the stage is processing a predetermined token (known in this example as the DATA token), then it will duplicate... the address field of the DATA token. If, on the other hand, the stage is processing any other kind of token, it will delete every word. The overall effect is that... respective output signals:

In the duplication stage, the output from the data latch LDIN forms intermediate data referred to as MID(underscore)DATA. This intermediate data word is loaded into the data output latch LDOUT only when an intermediate acceptance signal (labeled "MID(underscore)ACCEPT" in Fig. 8a) is set HIGH.

The portion of... data. These include a "DATA(underscore)TOKEN" signal that indicates that the circuitry is currently processing a valid DATA Token, and a NOT(underscore)DUPLICATE signal which is used to control duplication of data. When the circuitry is processing a DATA Token, the NOT(underscore)DUPLICATE signal toggles between a HIGH and a LOW... the token to be duplicated once (but no more times). When the circuitry is not processing a valid DATA Token then the NOT(underscore)DUPLICATE signal is held in a HIGH state. Accordingly, this means that the token words that are being processed are not duplicated.

As Fig. 8a illustrates, the upper six bits of 8-bit intermediate data word and the output signal QI1 from the latch LI1 form inputs to a... explained further below.

Latch LO1 performs the function of latching the last value of the intermediate extension bit (labeled "MID(underscore)EXTN" and as signal S4), and it loads this value... and the DATA(underscore)TOKEN signal will become "0", indicating that the circuitry is not processing a DATA token.

If QI1 is "0" and SO is "0", thereby indicating a DATA... phase and the DATA(underscore)TOKEN signal will become "1", indicating that the circuitry is processing a DATA token.

The NOT(underscore)DUPLICATE signal (the output signal QO3) is similarly loaded...LVOUT at the same time that MID(underscore)DATA is loaded into LDOUT and the intermediate extension bit (signal S4) is loaded into LEOUT. Signal S5 is also combined with the...above. This has the effect that all tokens except the one that causes the duplication process will be deleted from the token stream, since a device connected to the output terminals (OUTDATA, OUTEXTN and OUTVALID) will not recognize these token words as valid data.

As before... and is duplicated.

Referring now more particularly to Figure 10, there is shown a reconfigurable process stage in accordance with one aspect of the present invention.

Input latches 34 receive an... the input latches 34 is passed as a first input over line 35 to a processing unit 36. A first output from the token decode subsystem 33 is passed over line 37 as a second input to the

processing unit 36. A second output from the token decode 33 is passed over line 40 to an action identification unit 39. The action identification unit 39 also receives input from registers 43 and 44 over line 46. The registers 43... is determined by the history of tokens previously received. The output from the action identification unit 39 is passed over line 38 as a third input to the processing unit 36. The output from the processing unit 36 is passed to output latches 41. The output from the output latches 41 is... decoder 56 is passed over line 63 as an input to an Index to Data Unit (ITOD) 64. The Huffman decoder 56 and the ITOD 64 work together as a single logical unit.

The output from the ITOD 64 is passed over line 65 to an arithmetic logic unit (ALU) 66. A first output from the ALU 66 is passed over line 67 to...blocks 133.

Referring to Figure 14b, in the JPEG and H.261 standards, the Common Intermediate Format (CIF) is used, wherein a picture 141 is encoded as 6 rows each containing... ..in a zigzag direction indicated by the arrow 144. The GOBs 142 are, in turn, processed row-by-row, left-to-right in each row.

Referring now to Figure 14c, it... ..in accordance with the practice of the present invention. A first picture 161 to be processed contains a first PICTURE(underscore)START token 162, first-picture information of indeterminate length 163, and a first PICTURE(underscore)END token 164. A second picture 165 to be processed contains a second PICTURE(underscore)START token 166, second picture information of indeterminate length 167... ..tokens 162 and 166 indicate the start of the pictures 161 and 165 to the processor. Likewise, the PICTURE(underscore)END tokens 164 and 168 signify the end of the pictures 161 and 165 to the processor. This allows the processor to process picture information 163 and 167 of variable lengths.

Referring to Figure 17, a split 171...Video Formatter (not shown in Figure 17).

Referring now to Figure 18, the prediction filtering process is illustrated. A forward picture 201 is passed over line 202 as a first input... ..the right of the value decode shift register 230, as indicated by area 231. This process eliminates overlapping start code images, as discussed below. A first output from the value decode... ..Code Detector. The Start Code Detector then receives a first data value image 244. Before processing the first data value image 244, the Start Code Detector may detect a second start... ..image 244 at a length 246. If this occurs, the Start Code Detector does not process the first data value image 244, and instead receives and processes a second data value image 247.

...line 1 of Table 600, whenever a "sequence start" image is received during H.261 processing or a "picture start" image is received during MPEG processing, the entire group of four control tokens is generated, each followed by its corresponding data...Picture Decoding

3. Motion Picture Decompression

4. RAM Memory Map

5. Bitstream Characteristics

6. Reconfigurable Processing Stage

7. Multi-Standard Coding

8. Multi-Standard Processing Circuit-2nd Mode of Operation

9. Start Code Detector

10. Tokens

11. DRAM Interface

12... ..described herein in greater detail) and reformatting this output for use, including display in a computer or other display systems, including a video display system. Implementation of this formatting varies significantly... ..the Spatial Decoder circuits.

The Spatial Decoder of the present invention performs all the required processing within a single picture. This reduces the redundancy within one picture.

The Temporal Decoder reduces... ..modeller 75, the inverse zig-zag 81 and the inverse DCT 83. The standard independent units within the Huffman decoder and parser include the ALU 66 and the token formatter 71.

Referring now to Figure 12, the standard-independent units include the DRAM interface 100, the fork 91, the FIFO register 96, the summer 98 and the output selector 106. The standard dependent units are the address generator 94, which is different in H.261 and in MPEG, and...much of the operation is very similar between the three different compression standards.

The next unit is the state machine 68 (Figure 11) located within the Huffman decoder and parser. Here... The same holds true for JPEG, which is a third, completely independent program.

The next unit is the Huffman decoder 56 which functions with the index to data unit 64. Those two units cooperate together to perform the Huffman decoding. Here, the algorithm that is used for Huffman... to the Huffman decoder at different times consistent with the standard in operation.

The last unit on the chip that is dependent on the compression standard is the inverse quantizer 79...an H.261 group of blocks and an MPEG slice. When H.261 data is processed after the Start Code Detector, each group of blocks is preceded by a slice(underscore... these standards have totally different sets of tables.

As previously indicated, most of the system units are compression standard independent. If a unit is standard independent, and such units need not remember what CODING(underscore)STANDARD is being processed. All of the units that are standard dependent remember the compression standard as the CODING(underscore)STANDARD token flows... CODING(underscore)STANDARD tokens at the Start Code Detector that is positioned as the first unit in the pipeline, this change of compression standard is readily handled. The token says a... found in the standard, i.e. from the bitstream into a prediction mode token. This processing is performed by the Huffman decoder and parser state machine, where it is easy to... to that token. By having these tokens and using them appropriately, the design of other units in the machine is simplified. Although there may be some complications in the program, benefits... a first encoded signal (the MPEG or H.261 encoded video signal) in a pipeline processing system. The Temporal Decoder is not needed for JPEG decoding.

In this regard, the invention... the use of a single pipeline decoder and decompression system. The decoding and decompression pipeline processor is organized on a unique and special configuration which allows the handling of the multi... video signals through the use of techniques all compatible with the single pipeline decoder and processing system. The Spatial Decoder is combined with the Temporal Decoder, and the Video Formatter is...with only still pictures. The compression standard independent Spatial Decoder performs all of the data processing within the boundaries of a single picture. Such a decoder handles the spatial decompression of... to the multi-standard, configurable Video Formatter, which then provides an output to the display terminal. In a first sequence of similar pictures, each decompressed picture at the output of the... of control tokens and DATA tokens, in combination with a plurality of sequentially-positioned reconfigurable processing stages selected and organized to act as a standard-independent, reconfigurable-pipeline-processor.

With regard to JPEG decoding, a single Spatial Decoder with no off chip DRAM can... video. Accordingly, signals carried by DATA tokens pass directly through the Temporal Decoder without further processing when the Temporal Decoder is configured for a JPEG operation.

Another aspect of the present... for subsequent use in temporal decoding of subsequent pictures.

Generally, the Temporal Decoder performs the processing between pictures either earlier and/or later in time with reference to the picture currently... is distributed among several areas of DRAM in the sense that the decompressed output information, processed by the Spatial Decoder, is stored in other DRAM registers by other random access memories...first decoder circuit (the Spatial Decoder) directly to the Video Formatter for handling without signal processing delay.

The Temporal Decoder also reorders the blocks of picture data for display by a... from a selection of pictures which have arrived earlier or later than the picture under processing. When a picture is described in this context, it may

mean any one of the... 2. The result, i.e., the final decoded picture resulting from the addition of a process step performed by the decoder;

3. Previously decoded pictures read from the DRAM; and

4... START token and a subsequent PICTURE(underscore)END token.

After the picture data information is processed by the Temporal Decoder, it is either displayed or written back into a picture memory location. This information is then kept for further reference to be used in processing another different coded data picture.

Re-ordering of the MPEG encoded pictures for visual display...used to encode a referenced picture of a picture might be identified as being one unit long, another picture might be a number of units long, while still a third picture could be a fraction of that unit.

None of the existing standards (MPEG 1.2, JPEG, H.261) define a way of... picture rate, whereas the Video Formatter can handle a variable input picture rate.

6. RECONFIGURABLE PROCESSING STAGE

Referring again to Figure 10, the reconfigurable processing stage (RPS) comprises a token decode circuit 33 which is employed to receive the tokens... input latches 34. The output of the token decode circuit 33 is applied to a processing unit 36 over the two-wire interface 37 and an action identification circuit 39. The processing unit 36 is suitable for processing data under the control of the action identification circuit 39. After the processing is completed, the processing unit 36 connects such completed signals to the output, two-wire interface bus 40 through output... token decode circuit 33 are applied simultaneously to the action identification circuit 39 and the processing unit 36. The action identification function as well as the RPS is described in further detail... not standard independent circuits. The data flows through the token decode circuit 33, through the

processing unit 36 and onto the two-wire interface circuit 42 through the output latches 41. If... wire interface 42 through the output circuit 41. The present invention operates as a pipeline processor having a two-wire interface for controlling the movement of control tokens through the pipeline... time, the token decode circuit 33 provides a proper flag or index signal to the processing unit 36 to alert it to the presence of the token being handled by the action identification circuit 39.

Control tokens may also be processed.

A more detailed description of the various types of tokens usable in the present invention... standard now passing through the state machine shown with reference to Figure 10.

Similarly, the processing unit 36 which is under the control of the action identification circuit 39 is now ready to process the information contained in the data fields of the DATA token when it is appropriate... action identification circuit 39 and is immediately followed by a DATA token which is then processed by the processing unit 36. The control token exits the output latches circuit 41 over the output two-wire interface 42 immediately preceding the DATA token which has been processed within the processing unit 36.

In the present invention, the action identification circuit, 39, is a state machine holding... show that the action can also be affected by the token that is currently being processed by the token decode circuit 33.

In general, there is shown token decoding and data processing in accordance with the present invention. The data processing is performed as configured by the action identification circuit 39. The action is affected by... information stored from previously decoded tokens in registers 43 and 44, the current token under processing, and

the state and history information that the action identification unit 39 has itself acquired. A distinction is thereby shown between Control tokens and DATA tokens.

In any RPS, some tokens are viewed by that RPS unit as being Control tokens in that they affect the operation of the RPS presumably at... are viewed by the RPS as DATA tokens. Such DATA tokens contain information which is processed by the RPS in a way that is determined by the design of the particular... view of the same token. Some of the tokens might be viewed by one RPS unit as DATA Tokens while another RPS unit might decide that it is actually a Control Token. For example, the quantization table information... into a token called a quantization table token (QUANT(TABLE)) which goes down the processing pipeline. As far as that machine is concerned, all of that was data; it was... sort of data into another sort of data, which is clearly a function of the processing performed by that portion of the machine. However, when that information gets to the inverse... present. This information is viewed as control information, and then that control information affects the processing that is done on subsequent DATA tokens because it affects the number that you multiply... important feature of the invention is that each of the stages of circuitry has the processing capability within it to be able to perform the necessary operations for each of the... operations are to be performed at a given time, come as tokens. There is one processing element that differs between the different stages to provide this capability. In the state machine... standard is and it looks up the parameters that it needs to apply to the processing elements in order to perform a proper operation. For example, the inverse quantizer will look... is set to 1 for a particular compression standard, and will apply that to its processing circuitry.

In a similar sense the Huffman decoder 56 has a number of tables within... MPEG video standard or the JPEG video standard. These three compression coding standards specify similar processes to be done on the arriving data, but the structure of the datastreams is different... token stream embodying the current coding standard. The control tokens are passed through the pipeline processor, and are used, i.e., decoded, in the state machines to which they are relevant... this regard, the DATA Tokens are treated in the same fashion, insofar as they are processed only in the state machines that are configurable by the control tokens into processing such DATA Tokens. In the remaining state machines, they pass through unchanged.

More specifically, a... signals. The remaining portions of the token are used to indicate and identify the internal processing control function which is standard for all of the datastreams passing through the pipeline processor. In one form of the invention, the token extension is used to carry the current... accompanying data. As previously discussed, this information is utilized in the system to reconfigure the processing stage used to perform the function required by the various standards created for that purpose... picture number as indicated by the value.

The system also includes a multi-stage parallel processing pipeline operating under the principles of the two-wire interface previously described. Each of the... the token presently entering the state machine into the action identification circuit 39 or the processing unit 36, as appropriate. The processing unit has been previously reconfigured by the next previous control token into the form needed for handling the current coding standard, which is now entering the processing stage and carried by the next DATA token. Further, in accordance with this aspect of the invention, the succeeding state machines in the processing pipeline can be functioning under one coding standard, i.e., H.261, while a previous... tokens required to decode a number of coding standards with a fixed number of reconfigurable processing stages. More specifically, the PICTURE(END) control token is employed because it is important... standard machine, it is necessary to create additional control tokens within the multi-standard pipeline processing machine which will then indicate which one of the standard decoding techniques to use. Such... and to push the current picture through the decoder to the display.

8. MULTI-STANDARD PROCESSING CIRCUIT - SECOND MODE OF OPERATION

A compression standard-dependent circuit, in the form of the... of the Start Code Detector will subsequently be discussed in further detail, as will the process of starting up of the decoder.

The aforementioned description has been concerned primarily ...the data which immediately follows according to the standard. However, in the multi-standard pipeline processing system of the present invention, where compatibility is required for multiple standards, the system has... signals, including flag signals, are generated by each state machine to handle some of the processing within that state machine. Values carried in the standards can be used to access machine... its contents must be removed from the two wire interface to ensure that no further processing takes place using these 3 bytes. The decode register is emptied, and the value decode... 10. TOKENS

In the practice of the present invention, a token is a universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions and is adapted for use with a reconfigurable processing stage (RPS) which is a stage, which in response to a recognized token, reconfigures itself to perform various operations.

Tokens may be either position dependent or position independent upon the processing stages for performance of various functions. Tokens may also be metamorphic in that they can be altered by a processing stage and then passed down the pipeline for performance of further functions. Tokens may interact... other functions, and the specific interaction with a stage may be conditioned by the previous processing history of a stage.

A PICTURE(underscore)END token is a way of signalling the... through a fixed size, fixed width buffer.

The present invention is directed to a pipeline processing system which has a variable configuration which uses tokens and a two-wire system. The... do not use control tokens.

The control tokens are generated by circuitry within the decoder processor and emulate the operation of a number of different type standard-dependent signals passing into the serial pipeline processor for handling. The technique used is to study all the parameters of the multi-standards that are selected for processing by the serial processor and noting 1) their similarities, 2) their dissimilarities, 3) their needs and requirements and 4) selecting the correct token function to effectively process all of the standard signals sent into the serial processor. The functions of the tokens are to emulate the standards. A control token function is... the standard dependent signals and as an element to transmit control information through the pipeline processor.

In prior art system, a dedicated machine is designed according to well-known techniques to... tokens provide and make a sensible format for communicating information through the decompression circuit pipeline processor. In the design selected hereinafter and used in the preferred embodiment, each word of a... However, this is not a limitation on the invention, but on the magnitude of the processing steps elected to be accomplished by use of these tokens. It is to be noted... bit address for use in accessing the random access memories used throughout this serial decompression processor. This provides an additional degree of variability that facilitates a broad range of versatility.

As previously described, the DATA token carries data from one processing stage to the next. Consequently, the characteristics of this token change as it passes through... longest number of data bits because it needs to provide the most information to the

processing unit so that it can start the decompression with as much information as possible. Words which...to receive an address, it waits for the address generator to supply a valid address, processes that address and then sets the accept line high for one clock period. Thus, it... be read. This signal passes between two asynchronous clock regimes and, therefore, passes through three synchronizing flip flops.

Provided RAM2 312 is empty, the next item of data to arrive on...interesting.

In general, prediction data will be offset from the position of the block being processed as specified in the motion vectors in x and y. Thus, the block of data... address, 9. Data is read from this address and the x value is incremented. The process is repeated until the x value reaches its stop value, at which point, the y... is read, the x value is again incremented until it reaches its stop value. The process is repeated until both x and y values have

reached their stop values. Thus, the...invention, is that additional information must be provided to the prediction filters to indicate what processing is required on the data. This consists of the following:

a "last byte" signal indicating... ..bit 0) is incremented and the x address (3 LSBS) is reset to zero. This process is repeated until 64 bytes have been read. With a 16 or 32 bit wide...register while its access register is set to zero, the results are undefined.

14. MICRO-PROCESSOR INTERFACE

A standard byte wide micro-processor interface (MPI) is used on all circuits with in the Spatial Decoder and Temporal Decoder... ..the parameter column. The actual specifications are shown in the respective columns min, max and units.

The DC operating conditions can be seen with reference to Table A.6.3. Here... ..the signal is present the maximum amount of time that this signal is available. The Units column gives the units of measurement used to describe the signals.

16. MPI WRITE TIMING

The general description of...a PICTURE(underscore)END token is decoded and forces the data in the coded data buffers to be applied to the Huffman decoder and video demultiplexor, the final picture can be... ..Consequently, the machine will not go into error recovery mode and will successfully continue to process the coded data.

A still further advantage of the use of a PICTURE(underscore)END token is that the serial pipeline processor will continue the processing of uninterrupted data. Through the use of a PICTURE(underscore)END token, the serial pipeline processor is configured to handle less than the expected amount of data and, therefore, continues processing. Typically, a prior art machine would stop itself because of an error condition. As previously... ..of the Huffman decode and Video Demultiplexor know the number of blocks that it will process during each picture recovery cycle. When the correct number of blocks do not arrive from... ..Each of the state machines recognizes a FLUSH control token as information not to be processed. Accordingly, the FLUSH token is used to fill up all of the remaining empty parts... ..Huffman Decoder and Video Demultiplexor. In this way, the FLUSH token is like padding for buffers.

The Token Decoder in the Huffman circuit recognizes the FLUSH token and ignores the pseudo... ..less information than normally expected to decode the last picture. The Huffman decode circuit finishes processing the information contained in the last picture, and outputs this information through the DRAM interface... ..token, in accordance with the present invention, is used to pass through the entire pipeline processor and to ensure that the buffers are emptied and that other circuits are reconfigured to... ..underscore)END token, a padding word and a FLUSH token indicating to the serial pipeline processor that the picture processing for the current picture form is completed. Thereafter, the various state machines need reconfiguring to... ..FLUSH token resets each stage as it passes through, but-allows subsequent stages to continue processing. This prevents a loss of data. In other words, the FLUSH token is a variable... ..AFTER PICTURE

The STOP(underscore)AFTER(underscore)PICTURE function is employed to shut down the processing of the serial pipeline decompressing circuit at a logical point in its operation. At this... ..a picture, the STOP(underscore)AFTER(underscore)PICTURE operation signals the end of all current processing.

22. MULTI(underscore)STANDARD - SEARCH MODE

Another feature of the present invention is the use... ..underscore)MODE control token which is used to reconfigure the input to the serial pipeline processor to look at the incoming bit stream. When the search mode is set, the Start... ..combination of control tokens, and DATA tokens along with the reconfiguration circuits, to provide similar processing.

The use of search mode in the present invention is convenient in many situations including... video disc. In general, a search mode is convenient when the user interrupts the normal processing of the serial pipeline at a point where the machine does not expect such an... be the case.

In brief, the Huffman Decoder 321 works in conjunction with the other units shown in Figure 27. These other units are the Parser State Machine 322, the inshifter 323, the Index to Data unit 324, the ALU 325, and the Token Formatter 326. As described previously, connection between these blocks is governed by a two wire interface. A more detailed description of how these units function is subsequently described herein in greater detail, the focus here is on particular aspects... control certain functions of the Index to Data 324 and ALU 325. Control of these units by the Huffman Decoder is necessary for proper decoding of block-level information. Having the... further detail in the "More Detailed Description of the Invention" section.

The Index to Data unit 324 performs the second part of the multi-part algorithm. This unit contains a look up table that provides the actual Huffman decoded data. Entries in the... by detecting these in the Huffman Decoder 321, rather than in the Index to Data unit 324.

This index number is then passed to the Index to Data unit 324. In essence, the Index to Data unit is a look-up table. In accordance with one aspect of the algorithm, the look... format that JPEG specifies for transferring an alternate JPEG table.

From the Index to Data unit 324, the decoded index number or other data is passed, together with the accompanying control... the entering data to ensure that the DATA tokens are of the correct size for processing. In fact, the token stream can be corrected in some situations if the error is... an order that is useful for the decompression circuits, but not for the particular display unit being used. When a block of data enters the Buffer Manager, the Buffer Manager supplies... the output of the Spatial Decoder or Temporal Decoder and re-format it for a computer or display system. The details of this formatting will vary between applications. In a simple... Token. The DATA Token can have as many bits as are necessary for carrying out processing at a particular place in the system. All other Tokens ignore the extra bits.

A.3.2 The DATA Token

The DATA Token carries data from one processing stage to the next. Consequently, the characteristics of this Token change as it passes through... will be sufficient to collect DATA Tokens and to detect a few Tokens that provide synchronization information (such as PICTURE(underscore)START). In this regard, see subsequent sections A.16, "Connecting... from the data stream. This provides an alternative to doing the configuration via the micro processor interface.

A.3.4 Description of Tokens

This section documents the Tokens which are implemented... 3.5.1. Note: JPEG requires a 2:1:1 structure for its macroblocks when processing 4:2:2 data. See Table A.3.5.

A.3.6 Special Token formats... either is low then the interface is taken to high impedance.

Note: on-chip data processing is not terminated when the DRAM interface is at high impedance. Therefore, errors will occur... decoded video's picture rate. Accordingly, this clock can be used to provide audio/video synchronization.

A.7.1 Spatial Decoder clock signals

The Spatial Decoder has two different (and potentially... in accordance with the present invention, must know what video standard is being input for processing. Thereafter, the system can accept either pre-existing Tokens or raw byte data which is... time a value is written into coded(underscore)data (7:0). Software is responsible for settling

coded(underscore)extn to 0 before the last word of any Token is written to... ..0). The start of this new DATA Token then passes into the Spatial Decoder for processing.

Each time a new 8 bit value is written to coded(underscore)data (7:0... ..Detector analyses data in the DATA Tokens bit serially. The Detector's normal rate of processing is one bit per clock cycle (of coded(underscore)clock). Accordingly, it will typically decode a byte of coded data every 8 cycles of coded(underscore)clock. However, extra processing cycles are occasionally required, e.g., when a non-DATA Token is supplied or when the main decoder(underscore)clock. Data transfer is synchronized to decoder(underscore)clock on-chip.

SECTION A.11 Start code detector

A.11.1... ..Code Detector. So, accessing these registers will be unreliable if the Start Code Detector is processing data. The user is responsible for ensuring that the Start Code Detector is halted before... ..Detector. In this case, the Tokens are passed through the Start Code Detector with no processing to other stages of the Spatial Decoder. These Tokens can only be inserted just before...result will be unpredictable if this is done when the Start Code Detector is actively processing data.

Discard all mode can be safely initiated after any of the Start Code Detector...start code non-alignment interrupt is suppressed.

In contrast, however, JPEG was designed for a computer environment where byte alignment is guaranteed. Therefore, marker codes should only be detected when byte... ..the other hand, was designed to meet the needs of both communications (bit serial) and computer (byte oriented) systems. Start codes in MPEG data should normally be byte aligned. However, the...result will be unpredictable if this is done when the Start Code Detector is actively processing data. So, before initiating a start code search, the Start Code Detector should be stopped so no data is being processed. The Start Code Detector is always in this condition if any of the Start Code...the spatial video decoding circuits (inverse modeler, quantizer and DCT). This second logical buffer allows processing time to include a spread so as to accommodate processing pictures having varying amounts of data.

Both buffers are physically held in a single off 1.1, the unit for all the above mentioned registers is a 512 bit block of data. Accordingly, the... ..until there is space in the buffer. If a buffer continues to be full, more processing stages "up steam" of the buffer will halt until the Spatial Decoder is unable to... ..converting coded data into Tokens started by the Start Code Detector. There are four main processing blocks in the Video Demux: Parser State Machine, Huffman decoder (including an ITOD), Macroblock counter... ..or state machine follows the syntax of the coded video data and instructs the other units. The Huffman decoder converts variable length coded (VLC) data into integers. The Macroblock counter keeps...

Claims: ...a video decoding and decompression system having an input, an output and a plurality of processing stages between the input and the output defining a pipeline, the improvement comprising :

a token... ..received via said input for generating an interactive interfacing control token, defining a universal adaptation unit, for data functions among said processing stages, wherein said token is variable in length and is transmitted serially through said processing stages of said pipeline, and wherein said token is altered by a said processing stage ;

a first two wire interface disposed between a preceding member and a succeeding member... ..to enable loading of data and validation signals into said respective storage devices ;

whereby said processing stages are afforded enhanced flexibility in the processing of data.

2. The system as recited in claim 1, wherein said token is generated by one of said processing stages.

3. The system as recited in any one of claims 1 and 2, wherein... ..stages.

4. The system as recited in claim 1, wherein said control token causes said processing stages to reconfigure.
5. The system as recited in claim 1, wherein the interaction of said token with a stage is conditioned by the previous processing history of said stage.
6. The system as recited in claim 1, wherein said token... ..said token.
7. A system as recited in claim 6, wherein interaction with a selected processing stage is determined by said address field.
8. The system as recited in claim 1... ..being determined by said extension bits.
9. The system according to claim 1, wherein said processing stages comprise :
a temporal decoder ;
a spatial decoder ; and
a video formatter ; the system further... ..wherein said memory interface is clocked asynchronous with said address generator and with another said processing stage that provides the data being transmitted through said memory interface.
11. The system according...

6/K/10 (Item 10 from file: 348)
EUROPEAN PATENTS
(c) 2008 European Patent Office. All rights reserved.

Start code detecting apparatus for a video data stream

Country	Number	Kind	Date
---------	--------	------	------

Abstract ...said extension indicators, whereby the length of said token can be unlimited ;

Type	Pub. Date	Kind	Text
Available Text	Language	Update	Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: INTRODUCTION

The present invention is directed to improvements in methods and apparatus for decompression which operates to decompress and/or decode a plurality of differently encoded input... ..of the well known standards known as JPEG, MPEG and H.261.

A serial pipeline processing system of the present invention comprises a single two-wire bus used for carrying unique... ..to a plurality of adaptive decompression circuits and the like positioned as a reconfigurable pipeline processor.

PRIOR ART

One prior art system is described in United States Patent No. 5,216,724. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and a host processor. The device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus.

United States Patent No. 4,785... ...a known quad tree data structure.

United States Patent No. 5,122,875 discloses an apparatus for encoding/decoding an HDTV signal. The apparatus includes a compression circuit responsive to high definition video source signals for providing hierarchically layered... ...compressed video data of relatively greater and lesser importance to image reproduction respectively. A transport processor, responsive to the high and low priority codeword sequences, forms high and low priority transport... ...United States Patent No. 5,168,356 discloses a video signal encoding system that includes apparatus for segmenting encoded video data into transport blocks for signal transmission. ...in respective transport blocks.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening. This is accomplished by an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ...frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ...oscillators and the receiver can continuously receive each channel, then the receiver need not be synchronized with the transmitter. An FFT algorithm implements a fast discrete approximation to the continuous case in which the receiver synchronizes to the first frame and then acquires subsequent frames every frame period. The frame period... ...increasing the amount of data transmitted.

United States Patent No. 5,212,742 discloses an apparatus and method for processing video data for compression/decompression in real-time. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and host processor. Lastly, the device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus. The method handles assigning portions of the image for each of the processors to operate upon.

United States Patent No. 5,231,484 discloses a system and method... ...MPEG standards. Included are three cooperating components or subsystems that operate to variously adaptively pre-process the incoming digital motion video sequences, allocate bits to the pictures in a sequence, and... ...States Patent No. 5,267,334 discloses a method of removing frame redundancy in a computer system for a sequence of moving images. The method comprises detecting a first scene change... ...facing" keyframe or intraframe, and it is normally present in CCITT compressed video data. The process then comprises generating at least one intermediate compressed frame, the at least one intermediate compressed frame containing difference information from the first image for at least one image following...change, known as a "backward-facing" keyframe. The first keyframe and the at least one intermediate compressed frame are linked for forward play, and the second keyframe and the intermediate compressed frames are linked in reverse for reverse play. The intraframe may also be used... ...of complete scene information.

United States Patent No. 5,276,513 discloses a first circuit apparatus, comprising a given number of prior-art image-pyramid stages, together with a second circuit apparatus, comprising the same given number of novel motion-vector stages, perform cost-effective hierarchical motion analysis (HMA) in real-time, with minimum system processing delay and/or employing minimum system processing delay and/or employing minimum hardware structure. Specifically, the first and second circuit apparatus, in response to relatively high-resolution image data from an ongoing input series of successive... ...a relatively high frame rate (e.g., 30 frames per second), derives, after

a certain processing-system delay, an ongoing output series of successive given pixel-density vector-data frames that... ..of successive image frames.

United States Patent No. 5,283,646 discloses a method and apparatus for enabling a real-time video encoding system to accurately deliver the desired number of... ..desired bit allocations.

The article, Chong, Yong M., A Data-Flow Architecture for Digital Image Processing, Wescon Technical Papers: No. 2 Oct./Nov. 1984, discloses a real-time signal processing system specifically designed for image processing. More particularly, a token based data-flow architecture is disclosed wherein the tokens are of... ..width having a fixed width address field. The system contains a plurality of identical flow processors connected in a ring fashion. The tokens contain a data field, a control field and a tag. The tag field of the token is further broken down into a processor address field and an identifier field. The processor address field is used to direct the tokens to the correct data-flow processor, and the identifier field is used to label the data such that the data-flow processor knows what to do with the data. In this way, the identifier field acts as an instruction for the data-flow processor. The system directs each token to a specific data-flow processor using a module number (MN). If the MN matches the MN of the particular stage... ..to locate the decoder in the preceding stage in order to pre-decode complex decoding processing and to alleviate critical path problems in the logic circuit. The elastic nature of the...of block signal in most cases.

United States Patent No. 4,903,018 discloses a process and data processing system for compressing and expanding structurally associated multiple data sequences. The process is particular to data sets in which an analysis is made of the structure in... ..data series on the basis of the order number of these data elements. The data processing system for performing the processes includes a storage matrix (26) and an index storage (28) having line addresses of the... ..the final actual video.

United States Patent No. 5,060,242 discloses an image signal processing system DPCM encodes the signal, then Huffman and run length encodes the signal to produce... ..tightly packed without gaps for efficient transmission without loss of any data. The tightly packed apparatus has a barrel shifter with its shift modulus controlled by an accumulator receiving code word... ..OR gate is connected to the shifter, while a register is connected to the gate. Apparatus for processing a tightly packed and decorrelated digital signal has a barrel shifter and accumulator for unpacking... ..an inverse DCPM decoder.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening is accomplished by use of an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ..frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ..kernel matrix.

United States Patent No. 5,231,486 discloses a high definition video system processes a bitstream including high and low priority variable length coded Data words. The coded Data... ..packed High Priority Data and packed Low Priority Data by means of respective data packing units. The coded Data is continuously applied to both packing units. High Priority and Low Priority Length words indicating the bit lengths of high priority and... ..States Patent No. 5,287,178 discloses a video signal encoding system includes a signal processor for segmenting encoded video data into transport blocks having a header section and a packed data section. The system also includes reset control apparatus for releasing resets of system components, after a global system reset, in a prescribed non-simultaneous phased sequence to enable signal processing to commence in the prescribed sequence. The phased reset release sequence begins when valid data...United States Patent No. 5,142,380 to Sakagami et al. discloses an image compression apparatus

suitable for use with still images such as those formed by electronic still cameras using... ..and Q.

United States Patent No. 5,193,002 to Guichard et al. disclosed an apparatus for coding/decoding image signals in real time in conjunction with the CCITT standard H.261. A digital signal processor carries out direct quantization and reverse quantization.

United States Patent No. 5,241,383 to Chen et al. describes an apparatus with a pseudo-constant bit rate video coding achieved by an adjustable quantization parameter. The... ..relates to an improved pipeline system having an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along the pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the processing stages in the pipeline and interacting with selected stages in the pipeline for control data and/or combined control-data functions among the processing stages, so that the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing. In accordance with the invention, the processing stages may be configurable in response to recognition of at least one token. One of the processing stages may be a Start Code Detector which receives the input and generates and/or... ..and resetting the system, and a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards. The present invention... ..data and having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video pictures of varying data size can be controlled. Also in accordance with the invention, a processing stage receives the input data stream, the stage including means for recognizing specified bit stream patterns, whereby the processing stage facilitates random access and error recovery. The invention may also include a means for... ..invention also includes an inverse modeller stage, an inverse discrete cosine transform stage, and a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for processing data.

In addition, the present invention relates to an improved pipeline system having a Huffman...pipeline stage that incorporates a two-wire transfer control and also shows two consecutive pipeline processing stages with the two-wire transfer control;

Figures. 5a and 5b taken together depict one... ..shown in Figures. 8a and 8b.

Figure 10 is a block diagram of a reconfigurable processing stage;

Figure 11 is a block diagram of a spatial decoder;

Figure 12 is a... ..decoder including the prediction filters;

Figure 18 is a pictorial representation of the prediction filtering process;

Figure 19 shows a generalized representation of the macroblock structure;

Figure 20 shows a generalized... ..buffer;

Figure 25 is a pictorial diagram illustrating prediction data offset from the block being processed;

Figure 26 is a pictorial diagram illustrating prediction data offset by (1,1);

Figure 27...in general terms, the present invention provides an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along a pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the stages in the pipeline and interacting with selected stages in the pipeline for control, data and/or combined control-data functions among the processing stages, whereby the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing.

Each of the processing stages in the pipeline may include both primary and secondary storage, and the stages in processing stages for performance of functions or position independent of the processing stages for performance of functions.

In a pipeline machine, in accordance with the invention, the... ..altered by interfacing with the stages, and the tokens may interact with all of the processing stages in the pipeline or only with some but less than all of said processing stages. The tokens in the pipeline may interact with adjacent processing stages or with non-adjacent processing stages, and the tokens may reconfigure the processing stages. Such tokens may be position dependent for some functions and position independent for other... ..be Huffman coded.

In the improved pipeline machine, the tokens may be generated by a processing stage. Such pipeline tokens may include data for transfer to the processing stages or the tokens may be devoid of data. Some of the tokens may be identified as DATA tokens and provide data to the processing stages in the pipeline, while other tokens are identified as control tokens and only condition the processing stages in the pipeline, such conditioning including reconfiguring of the processing stages. Still other tokens may provide both data and conditioning to the processing stages in the pipeline. Some of said tokens may identify coding standards to the processing stages in the pipeline, whereas other tokens may operate independent of any coding standard among the processing stages. The tokens may be capable of successive alteration by the processing stages in the pipeline.

In accordance with the invention, the interactive flexibility of the tokens in cooperation with the processing stages facilitates greater functional diversity of the processing stages for resident structure in the pipeline, and the flexibility of the tokens facilitates system... ..or alteration. The tokens may be capable of facilitating a plurality of functions within any processing stage in the pipeline. Such pipeline tokens may be either hardware based or software based... ..system bandwidth in the pipeline. The tokens may provide data and control simultaneously to the processing stages in the pipeline.

The invention may include a pipeline processing machine for handling plurality of separately encoded bit streams arranged as a single serial bit... ..and for passing unrecognized control tokens along the pipeline, and a reconfigurable decode and parser processing means responsive to a recognized control token for reconfiguring a particular stage to handle an... ..be a pipeline system and the Start Code Detector may be positioned as the first processing stage in the pipeline.

The present invention also provides, in a system having a plurality of processing stages, a universal adaptation unit in the form of an interactive interfacing token for control and/or data functions among the processing stages, the token being a PICTURE(underscore)START code token for indicating that the start... ..The token may also be a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards.

The CODING(underscore... ..standard as JPEG, and/or any other appropriate picture standard. At least some of the processing stages reconfigure in response to the CODING(underscore)STANDARD token.

One of the processing stages in the system may be a Huffman decoder and parser and, upon receipt of... ..Data stage, and the parser stage may send an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... ..video data, having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video... ..controlled.

The system may include a spatial decoder having a two-wire interface interconnecting processing stages, the interface enabling serial processing for data and parallel processing for control.

As previously indicated, the system may further include a ROM having separate stored... ..of a plurality of picture standards, the programs being selectable by a token to facilitate processing for a plurality of different picture standards.

The spatial decoder system also includes a token... ..decoding stage and a parser stage for sending an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... ..The present invention also provides a pipeline system having an input data stream, and a processing stage for receiving the input data stream, the stage including means for recognizing specified bit... ..whereby said stage facilitates random access and error recovery. In accordance with the invention, the processing stage may be a start code detector and the bit stream patterns may include start... ..token and padding insures uniformity of word size. In accordance with the invention, a reconfigurable processing stage may be provided as a spatial decoder and the padding means adds to picture... ..that if the DATA token has less than the predetermined length, the padder circuit adds units of data to the DATA token until the predetermined length is achieved. A bypass circuit...1 tokens into a buffer, having a second predetermined width.

The invention also provides an apparatus for providing a time delay to a group of compressed pictures, the pictures corresponding to... ..and capable of delaying the words of data, is in communication with a control circuit intermediate the counter circuit and the inverse modeller circuit, the control circuit also communicating with the... ..inverse modeller stage and an inverse discrete cosine transform stage, the improvement characterized by a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for

processing data.

In accordance with the invention, the token may be a QUANT(underscore)TABLE token for causing the processing stage to generate a quantization table.

The present invention also provides a Huffman decoder for... ..of bits used to represent an item of data.

DECODER: An embodiment of a decoding process.

DECODING (PROCESS): The process defined in this specification that reads an input coded bitstream and produces decoded pictures or... ..the same order in which they were presented at the input of the encoder.

ENCODING (PROCESS): A process, not specified in this specification, that reads a stream of input pictures or audio samples...to provide an estimate of the pel value or data element currently being decoded.

RECONFIGURABLE PROCESS STAGE (RPS): A stage, which in response to a recognized token, reconfigures itself to perform various operations.

SLICE: A series of macroblocks.

TOKEN: A universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions... ..indicates that the corresponding stage holds valid data, i.e., data that is to be processed in one of the pipeline stages. After processing (which may involve nothing more than a simple transfer without manipulation of the data) valid... ..present invention may be used with any number of pipeline stages. Furthermore, data may be processed in more than one stage and the processing time for different stages can differ.

In addition to clock and data signals (described below... ..other system. For example, the last pipeline stage may pass its data on to subsequent processing circuitry. The ACCEPT signal, which is illustrated as the lower of the two lines connecting... ..the minimum disturbance possible to other pipeline stages. Succeeding pipeline stages are allowed to continue processing and, therefore, this means that gaps open up in the stream of data following the... ..The data in the pipeline is encoded such that many different types of data are processed in the pipeline. This encoding accommodates data packets of variable size and the size of...the other hand, it may generate itself, all or part of the data to be processed in the pipeline. Indeed, as is explained below, a "stage" may contain arbitrary

processing circuitry, including none at all (for simple passing of data) or entire systems (for example... values zero and 255 may not be used.

If such a picture were to be processed in a pipeline built in the practice of the present invention, then one of these... data must not be written over since it is data that must be saved for processing or use in a downstream device e.g., a pipeline stage, a device or a... connected to the pipeline upstream contains data D4 that is to be transferred into and processed in the pipeline. ...pipeline, in accordance with the preferred embodiments of the present invention, to "fill up" empty processing stages is highly advantageous since the processing stages in the pipeline thereby become decouple from one another. In other words, even though... data can be transferred into the pipeline and between stages even when one or more processing stages is blocked.

In the embodiment shown in Fig. 1, it is assumed that the... propagate all the way back to the beginning of the pipeline if there is some intermediate stage that is able to accept new data.

In the embodiment illustrated in Fig. 1... has been mentioned. It is to be further understood that each pipeline stage may also process the data it has received arbitrarily before passing it between its internal storage elements or... the portion of the pipeline that contains input and output storage elements and that arbitrarily processes data stored in its storage elements.

Furthermore, the "device" ... valid data, but also when a stage requires more than one clock phase to finish processing its data. This also can occur when it creates valid data in one or both... control the passage of data between adjacent storage elements. The VALID signal may also be processed in an analogous manner.

A great advantage of the two-wire interface (one wire for... In addition, two extra latches and a small number of gates are preferably added to process the ACCEPT and VALID signals that are associated with the data latches in each half... application so requires. The interface in accordance with this embodiment can also be used to process analog signals.

As discussed previously, while other conventional timing arrangements may be used, the interface... circuit B1, which may be provided to convert output data from input latch LDIN into intermediate data, which is then later loaded in an output data latch LDOUT, which comprises the... is connected either directly as an input to the validation output latch LVOUT, or via intermediate logic devices or circuits that may alter the signal.

Similarly, the output validation signal QVOUT... to the input of the validation input latch QVIN of the following stage, or via intermediate devices or logic circuits, which may alter the validation signal. This ...word.

Preferred Data Structure - "tokens"

In the sample application shown in Fig. 4, each stage processes all input data, since there is no control circuitry that excludes any stage from allowing... are connected together in a relatively simple configuration. The simplest configuration is a pipeline of processing steps. For example, in the one shown in Fig. 1. The use of tokens, however... flows from left to right in the diagram. Data enters the machine and passes into processing Stage A. This may or may not modify the data and it then passes the... advantage of the tokens is their ability to achieve this kind of communication. Since any processing stage that does not recognize a token simply passes it on unaltered to the next... is transmitted along with the address and data fields in each token so that a processing stage can pass on a token (which can be of arbitrary length) without having to... be the first word of a new token.

Note that although the simple pipeline of processing stages is particularly useful, it will be appreciated that tokens may be applied to more complicated configurations of processing elements. An example of a more complicated processing element is described below.

It is not necessary, in accordance with the present invention, to... has extension bits. An example of this is a token that activates a stage that processes video quantization values stored in a quantization table (typically a memory

device). For example, a...turn, is of great importance in video data pipeline systems since it ensures that all processing stages can be continuously running at full bandwidth.

In accordance to the present invention, in... some other chips in the set. This is advantageous both from the perspective of a customer and from that of a chip manufacturer. Even if modifications mean that all chips are...the end of a token (and hence the start of the next token) to be processed correctly (including simple non-manipulative transfer), even if the token is not recognized by the... block diagram of a pipeline stage whose function is as follows. If the stage is processing a predetermined token (known in this example as the DATA token), then it will duplicate... the address field of the DATA token. If, on the other hand, the stage is processing any other kind of token, it will delete every word. The overall effect is that... respective output signals:

In the duplication stage, the output from the data latch LDIN forms intermediate data referred to as MID(underscore)DATA. This intermediate data word is loaded into the data output latch LDOUT only when an intermediate acceptance signal (labeled "MID(underscore)ACCEPT" in Fig. 8a) is set HIGH.

The portion of... data. These include a "DATA(underscore)TOKEN" signal that indicates that the circuitry is currently processing a valid DATA Token, and a NOT(underscore)DUPLICATE signal which is used to control duplication of data. When the circuitry is processing a DATA Token, the NOT(underscore)DUPLICATE signal toggles between a HIGH and a LOW... the token to be duplicated once (but no more times). When the circuitry is not processing a valid DATA Token then the NOT(underscore)DUPLICATE signal is held in a HIGH state. Accordingly, this means that the token words that are being processed are not duplicated.

As Fig. 8a illustrates, the upper six bits of 8-bit intermediate data word and the output signal QI1 from the latch LI1 form inputs to a... explained further below.

Latch LO1 performs the function of latching the last value of the intermediate extension bit (labeled "MID(underscore)EXTN" and as signal S4), and it loads this value... and the DATA(underscore)TOKEN signal will become "0", indicating that the circuitry is not processing a DATA token.

If QI1 is "0" and SO is "0", thereby indicating a DATA... phase and the DATA(underscore)TOKEN signal will become "1", indicating that the circuitry is processing a DATA token.

The NOT(underscore)DUPLICATE signal (the output signal QO3) is similarly loaded...LVOUT at the same time that MID(underscore)DATA is loaded into LDOUT and the intermediate extension bit (signal S4) is loaded into LEOUT. Signal S5 is also combined with the...above. This has the effect that all tokens except the one that causes the duplication process will be deleted from the token stream, since a device connected to the output terminals (OUTDATA, OUTEXTN and OUTVALID) will not recognize these token words as valid data.

As before... and is duplicated.

Referring now more particularly to Figure 10, there is shown a reconfigurable process stage in accordance with one aspect of the present invention.

Input latches 34 receive an... the input latches 34 is passed as a first input over line 35 to a processing unit 36. A first output from the token decode subsystem 33 is passed over line 37 as a second input to the

processing unit 36. A second output from the token decode 33 is passed over line 40 to an action identification unit 39. The action identification unit 39 also receives input from registers 43 and 44 over line 46. The registers 43... is determined by the history of tokens previously received. The output from the action identification unit 39 is passed over line 38 as a third input to the processing unit 36. The output from the processing unit 36 is passed to output latches 41. The output from the output latches 41 is... decoder 56 is passed over line 63 as an input to an Index to Data Unit (ITOD) 64. The Huffman decoder 56 and the ITOD 64 work together as a single logical unit.

The output from the ITOD 64 is passed over line 65 to an arithmetic logic unit (ALU) 66. A first output from the ALU 66 is passed over line 67 to...blocks 133.

Referring to Figure 14b, in the JPEG and H.261 standards, the Common Intermediate Format (CIF) is used, wherein a picture 141 is encoded as 6 rows each containing... ..in a zigzag direction indicated by the arrow 144. The GOBs 142 are, in turn, processed row-by-row, left-to-right in each row.

Referring now to Figure 14c, it... ..in accordance with the practice of the present invention. A first picture 161 to be processed contains a first PICTURE(underscore)START token 162, first-picture information of indeterminate length 163, and a first PICTURE(underscore)END token 164. A second picture 165 to be processed contains a second PICTURE(underscore)START token 166, second picture information of indeterminate length 167... ..tokens 162 and 166 indicate the start of the pictures 161 and 165 to the processor. Likewise, the PICTURE(underscore)END tokens 164 and 168 signify the end of the pictures 161 and 165 to the processor. This allows the processor to process picture information 163 and 167 of variable lengths.

Referring to Figure 17, a split 171...Video Formatter (not shown in Figure 17).

Referring now to Figure 18, the prediction filtering process is illustrated. A forward picture 201 is passed over line 202 as a first input... ..the right of the value decode shift register 230, as indicated by area 231. This process eliminates overlapping start code images, as discussed below. A first output from the value decode... ..Code Detector. The Start Code Detector then receives a first data value image 244. Before processing the first data value image 244, the Start Code Detector may detect a second start... ..image 244 at a length 246. If this occurs, the Start Code Detector does not process the first data value image 244, and instead receives and processes a second data value image 247.

...line 1 of Table 600, whenever a "sequence start" image is received during H.261 processing or a "picture start" image is received during MPEG processing, the entire group of four control tokens is generated, each followed by its corresponding data...Picture Decoding

3. Motion Picture Decompression

4. RAM Memory Map

5. Bitstream Characteristics

6. Reconfigurable Processing Stage

7. Multi-Standard Coding

8. Multi-Standard Processing Circuit-2nd Mode of Operation

9. Start Code Detector

10. Tokens

11. DRAM Interface

12... ..described herein in greater detail) and reformatting this output for use, including display in a computer or other display systems, including a video display system. Implementation of this formatting varies significantly... ..the Spatial Decoder circuits.

The Spatial Decoder of the present invention performs all the required processing within a single picture. This reduces the redundancy within one picture.

The Temporal Decoder reduces... ..modeller 75, the inverse zig-zag 81 and the inverse DCT 83. The standard independent units within the Huffman decoder and parser include the ALU 66 and the token formatter 71.

Referring now to Figure 12, the standard-independent units include the DRAM interface 100, the fork 91, the FIFO register 96, the summer 98 and the output selector 106. The standard dependent units are the address generator 94, which is different in H.261 and in MPEG, and...much of the operation is very similar between the three different compression standards.

The next unit is the state machine 68 (Figure 11) located within the Huffman decoder and parser. Here... The same holds true for JPEG, which is a third, completely independent program.

The next unit is the Huffman decoder 56 which functions with the index to data unit 64. Those two units cooperate together to perform the Huffman decoding. Here, the algorithm that is used for Huffman... to the Huffman decoder at different times consistent with the standard in operation.

The last unit on the chip that is dependent on the compression standard is the inverse quantizer 79...an H.261 group of blocks and an MPEG slice. When H.261 data is processed after the Start Code Detector, each group of blocks is preceded by a slice(underscore... these standards have totally different sets of tables.

As previously indicated, most of the system units are compression standard independent. If a unit is standard independent, and such units need not remember what CODING(underscore)STANDARD is being processed. All of the units that are standard dependent remember the compression standard as the CODING(underscore)STANDARD token flows... CODING(underscore)STANDARD tokens at the Start Code Detector that is positioned as the first unit in the pipeline, this change of compression standard is readily handled. The token says a... found in the standard, i.e. from the bitstream into a prediction mode token. This processing is performed by the Huffman decoder and parser state machine, where it is easy to... to that token. By having these tokens and using them appropriately, the design of other units in the machine is simplified. Although there may be some complications in the program, benefits... a first encoded signal (the MPEG or H.261 encoded video signal) in a pipeline processing system. The Temporal Decoder is not needed for JPEG decoding.

In this regard, the invention... the use of a single pipeline decoder and decompression system. The decoding and decompression pipeline processor is organized on a unique and special configuration which allows the handling of the multi... video signals through the use of techniques all compatible with the single pipeline decoder and processing system. The Spatial Decoder is combined with the Temporal Decoder, and the Video Formatter is...with only still pictures. The compression standard independent Spatial Decoder performs all of the data processing within the boundaries of a single picture. Such a decoder handles the spatial decompression of... to the multi-standard, configurable Video Formatter, which then provides an output to the display terminal. In a first sequence of similar pictures, each decompressed picture at the output of the... of control tokens and DATA tokens, in combination with a plurality of sequentially-positioned reconfigurable processing stages selected and organized to act as a standard-independent, reconfigurable-pipeline-processor.

With regard to JPEG decoding, a single Spatial Decoder with no off chip DRAM can... video. Accordingly, signals carried by DATA tokens pass directly through the Temporal Decoder without further processing when the Temporal Decoder is configured for a JPEG operation.

Another aspect of the present... for subsequent use in temporal decoding of subsequent pictures.

Generally, the Temporal Decoder performs the processing between pictures either earlier and/or later in time with reference to the picture currently... is distributed among several areas of DRAM in the sense that the decompressed output information, processed by the Spatial Decoder, is stored in other DRAM registers by other random access memories...first decoder circuit (the Spatial Decoder) directly to the Video Formatter for handling without signal processing delay.

The Temporal Decoder also reorders the blocks of picture data for display by a... from a selection of pictures which have arrived earlier or later than the picture under processing. When a picture is described in this context, it may

mean any one of the... 2. The result, i.e., the final decoded picture resulting from the addition of a process step performed by the decoder;

3. Previously decoded pictures read from the DRAM; and

4... START token and a subsequent PICTURE(underscore)END token.

After the picture data information is processed by the Temporal Decoder, it is either displayed or written back into a picture memory location. This information is then kept for further reference to be used in processing another different coded data picture.

Re-ordering of the MPEG encoded pictures for visual display...used to encode a referenced picture of a picture might be identified as being one unit long, another picture might be a number of units long, while still a third picture could be a fraction of that unit.

None of the existing standards (MPEG 1.2, JPEG, H.261) define a way of... picture rate, whereas the Video Formatter can handle a variable input picture rate.

6. RECONFIGURABLE PROCESSING STAGE

Referring again to Figure 10, the reconfigurable processing stage (RPS) comprises a token decode circuit 33 which is employed to receive the tokens... input latches 34. The output of the token decode circuit 33 is applied to a processing unit 36 over the two-wire interface 37 and an action identification circuit 39. The processing unit 36 is suitable for processing data under the control of the action identification circuit 39. After the processing is completed, the processing unit 36 connects such completed signals to the output, two-wire interface bus 40 through output... token decode circuit 33 are applied simultaneously to the action identification circuit 39 and the processing unit 36. The action identification function as well as the RPS is described in further detail... not standard independent circuits. The data flows through the token decode circuit 33, through the

processing unit 36 and onto the two-wire interface circuit 42 through the output latches 41. If... wire interface 42 through the output circuit 41. The present invention operates as a pipeline processor having a two-wire interface for controlling the movement of control tokens through the pipeline... time, the token decode circuit 33 provides a proper flag or index signal to the processing unit 36 to alert it to the presence of the token being handled by the action identification circuit 39.

Control tokens may also be processed.

A more detailed description of the various types of tokens usable in the present invention... standard now passing through the state machine shown with reference to Figure 10.

Similarly, the processing unit 36 which is under the control of the action identification circuit 39 is now ready to process the information contained in the data fields of the DATA token when it is appropriate... action identification circuit 39 and is immediately followed by a DATA token which is then processed by the processing unit 36. The control token exits the output latches circuit 41 over the output two-wire interface 42 immediately preceding the DATA token which has been processed within the processing unit 36.

In the present invention, the action identification circuit, 39, is a state machine holding... show that the action can also be affected by the token that is currently being processed by the token decode circuit 33.

In general, there is shown token decoding and data processing in accordance with the present invention. The data processing is performed as configured by the action identification circuit 39. The action is affected by... information stored from previously decoded tokens in registers 43 and 44, the current token under processing, and

the state and history information that the action identification unit 39 has itself acquired. A distinction is thereby shown between Control tokens and DATA tokens.

In any RPS, some tokens are viewed by that RPS unit as being Control tokens in that they affect the operation of the RPS presumably at... are viewed by the RPS as DATA tokens. Such DATA tokens contain information which is processed by the RPS in a way that is determined by the design of the particular... view of the same token. Some of the tokens might be viewed by one RPS unit as DATA Tokens while another RPS unit might decide that it is actually a Control Token. For example, the quantization table information... into a token called a quantization table token (QUANT(TABLE)) which goes down the processing pipeline. As far as that machine is concerned, all of that was data; it was... sort of data into another sort of data, which is clearly a function of the processing performed by that portion of the machine. However, when that information gets to the inverse... present. This information is viewed as control information, and then that control information affects the processing that is done on subsequent DATA tokens because it affects the number that you multiply... important feature of the invention is that each of the stages of circuitry has the processing capability within it to be able to perform the necessary operations for each of the... operations are to be performed at a given time, come as tokens. There is one processing element that differs between the different stages to provide this capability. In the state machine... standard is and it looks up the parameters that it needs to apply to the processing elements in order to perform a proper operation. For example, the inverse quantizer will look... is set to 1 for a particular compression standard, and will apply that to its processing circuitry.

In a similar sense the Huffman decoder 56 has a number of tables within... MPEG video standard or the JPEG video standard. These three compression coding standards specify similar processes to be done on the arriving data, but the structure of the datastreams is different... token stream embodying the current coding standard. The control tokens are passed through the pipeline processor, and are used, i.e., decoded, in the state machines to which they are relevant... this regard, the DATA Tokens are treated in the same fashion, insofar as they are processed only in the state machines that are configurable by the control tokens into processing such DATA Tokens. In the remaining state machines, they pass through unchanged.

More specifically, a... signals. The remaining portions of the token are used to indicate and identify the internal processing control function which is standard for all of the datastreams passing through the pipeline processor. In one form of the invention, the token extension is used to carry the current... accompanying data. As previously discussed, this information is utilized in the system to reconfigure the processing stage used to perform the function required by the various standards created for that purpose... picture number as indicated by the value.

The system also includes a multi-stage parallel processing pipeline operating under the principles of the two-wire interface previously described. Each of the... the token presently entering the state machine into the action identification circuit 39 or the processing unit 36, as appropriate. The processing unit has been previously reconfigured by the next previous control token into the form needed for handling the current coding standard, which is now entering the processing stage and carried by the next DATA token. Further, in accordance with this aspect of the invention, the succeeding state machines in the processing pipeline can be functioning under one coding standard, i.e., H.261, while a previous... tokens required to decode a number of coding standards with a fixed number of reconfigurable processing stages. More specifically, the PICTURE(END) control token is employed because it is important... standard machine, it is necessary to create additional control tokens within the multi-standard pipeline processing machine which will then indicate which one of the standard decoding techniques to use. Such... and to push the current picture through the decoder to the display.

8. MULTI-STANDARD PROCESSING CIRCUIT - SECOND MODE OF OPERATION

A compression standard-dependent circuit, in the form of the... of the Start Code Detector will subsequently be discussed in further detail, as will the process of starting up of the decoder.

The aforementioned description has been concerned primarily ...the data which immediately follows according to the standard. However, in the multi-standard pipeline processing system of the present invention, where compatibility is required for multiple standards, the system has... ..signals, including flag signals, are generated by each state machine to handle some of the processing within that state machine. Values carried in the standards can be used to access machine... ..its contents must be removed from the two wire interface to ensure that no further processing takes place using these 3 bytes. The decode register is emptied, and the value decode... ..10. TOKENS

In the practice of the present invention, a token is a universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions and is adapted for use with a reconfigurable processing stage (RPS) which is a stage, which in response to a recognized token, reconfigures itself to perform various operations.

Tokens may be either position dependent or position independent upon the processing stages for performance of various functions. Tokens may also be metamorphic in that they can be altered by a processing stage and then passed down the pipeline for performance of further functions. Tokens may interact... ..other functions, and the specific interaction with a stage may be conditioned by the previous processing history of a stage.

A PICTURE(underscore)END token is a way of signalling the... ..through a fixed size, fixed width buffer.

The present invention is directed to a pipeline processing system which has a variable configuration which uses tokens and a two-wire system. The... ..do not use control tokens.

The control tokens are generated by circuitry within the decoder processor and emulate the operation of a number of different type standard-dependent signals passing into the serial pipeline processor for handling. The technique used is to study all the parameters of the multi-standards that are selected for processing by the serial processor and noting 1) their similarities, 2) their dissimilarities, 3) their needs and requirements and 4) selecting the correct token function to effectively process all of the standard signals sent into the serial processor. The functions of the tokens are to emulate the standards. A control token function is... ..the standard dependent signals and as an element to transmit control information through the pipeline processor.

In prior art system, a dedicated machine is designed according to well-known techniques to... ..tokens provide and make a sensible format for communicating information through the decompression circuit pipeline processor. In the design selected hereinafter and used in the preferred embodiment, each word of a... ..However, this is not a limitation on the invention, but on the magnitude of the processing steps elected to be accomplished by use of these tokens. It is to be noted... ..bit address for use in accessing the random access memories used throughout this serial decompression processor. This provides an additional degree of variability that facilitates a broad range of versatility.

As previously described, the DATA token carries data from one processing stage to the next. Consequently, the characteristics of this token change as it passes through... ..longest number of data bits because it needs to provide the most information to the

processing unit so that it can start the decompression with as much information as possible. Words which...to receive an address, it waits for the address generator to supply a valid address, processes that address and then sets the accept line high for one clock period. Thus, it... ..be read. This signal passes between two asynchronous clock regimes and, therefore, passes through three synchronizing flip flops.

Provided RAM2 312 is empty, the next item of data to arrive on...interesting.

In general, prediction data will be offset from the position of the block being processed as specified in the motion vectors in x and y. Thus, the block of data... ..address, 9. Data is read from this address and the x value is incremented. The process is repeated until the x value reaches its stop value, at which point, the y... ..is read, the x value is again incremented until it reaches its stop value. The process is repeated until both x and y values have

reached their stop values. Thus, the...invention, is that additional information must be provided to the prediction filters to indicate what processing is required on the data. This consists of the following:

a "last byte" signal indicating... ..bit 0) is incremented and the x address (3 LSBS) is reset to zero. This process is repeated until 64 bytes have been read. With a 16 or 32 bit wide...register while its access register is set to zero, the results are undefined.

14. MICRO-PROCESSOR INTERFACE

A standard byte wide micro-processor interface (MPI) is used on all circuits with in the Spatial Decoder and Temporal Decoder... ..the parameter column. The actual specifications are shown in the respective columns min, max and units.

The DC operating conditions can be seen with reference to Table A.6.3. Here... ..the signal is present the maximum amount of time that this signal is available. The Units column gives the units of measurement used to describe the signals.

16. MPI WRITE TIMING

The general description of...a PICTURE(underscore)END token is decoded and forces the data in the coded data buffers to be applied to the Huffman decoder and video demultiplexor, the final picture can be... ..Consequently, the machine will not go into error recovery mode and will successfully continue to process the coded data.

A still further advantage of the use of a PICTURE(underscore)END token is that the serial pipeline processor will continue the processing of uninterrupted data. Through the use of a PICTURE(underscore)END token, the serial pipeline processor is configured to handle less than the expected amount of data and, therefore, continues processing. Typically, a prior art machine would stop itself because of an error condition. As previously... ..of the Huffman decode and Video Demultiplexor know the number of blocks that it will process during each picture recovery cycle. When the correct number of blocks do not arrive from... ..Each of the state machines recognizes a FLUSH control token as information not to be processed. Accordingly, the FLUSH token is used to fill up all of the remaining empty parts... ..Huffman Decoder and Video Demultiplexor. In this way, the FLUSH token is like padding for buffers.

The Token Decoder in the Huffman circuit recognizes the FLUSH token and ignores the pseudo... ..less information than normally expected to decode the last picture. The Huffman decode circuit finishes processing the information contained in the last picture, and outputs this information through the DRAM interface... ..token, in accordance with the present invention, is used to pass through the entire pipeline processor and to ensure that the buffers are emptied and that other circuits are reconfigured to... ..underscore)END token, a padding word and a FLUSH token indicating to the serial pipeline processor that the picture processing for the current picture form is completed. Thereafter, the various state machines need reconfiguring to... ..FLUSH token resets each stage as it passes through, but-allows subsequent stages to continue processing. This prevents a loss of data. In other words, the FLUSH token is a variable... ..AFTER PICTURE

The STOP(underscore)AFTER(underscore)PICTURE function is employed to shut down the processing of the serial pipeline decompressing circuit at a logical point in its operation. At this... ..a picture, the STOP(underscore)AFTER(underscore)PICTURE operation signals the end of all current processing.

22. MULTI(underscore)STANDARD - SEARCH MODE

Another feature of the present invention is the use... ..underscore)MODE control token which is used to reconfigure the input to the serial pipeline processor to look at the incoming bit stream. When the search mode is set, the Start... ..combination of control tokens, and DATA tokens along with the reconfiguration circuits, to provide similar processing.

The use of search mode in the present invention is convenient in many situations including... video disc. In general, a search mode is convenient when the user interrupts the normal processing of the serial pipeline at a point where the machine does not expect such an... be the case.

In brief, the Huffman Decoder 321 works in conjunction with the other units shown in Figure 27. These other units are the Parser State Machine 322, the inshifter 323, the Index to Data unit 324, the ALU 325, and the Token Formatter 326. As described previously, connection between these blocks is governed by a two wire interface. A more detailed description of how these units function is subsequently described herein in greater detail, the focus here is on particular aspects... control certain functions of the Index to Data 324 and ALU 325. Control of these units by the Huffman Decoder is necessary for proper decoding of block-level information. Having the... further detail in the "More Detailed Description of the Invention" section.

The Index to Data unit 324 performs the second part of the multi-part algorithm. This unit contains a look up table that provides the actual Huffman decoded data. Entries in the... by detecting these in the Huffman Decoder 321, rather than in the Index to Data unit 324.

This index number is then passed to the Index to Data unit 324. In essence, the Index to Data unit is a look-up table. In accordance with one aspect of the algorithm, the look... format that JPEG specifies for transferring an alternate JPEG table.

From the Index to Data unit 324, the decoded index number or other data is passed, together with the accompanying control... the entering data to ensure that the DATA tokens are of the correct size for processing. In fact, the token stream can be corrected in some situations if the error is... an order that is useful for the decompression circuits, but not for the particular display unit being used. When a block of data enters the Buffer Manager, the Buffer Manager supplies... the output of the Spatial Decoder or Temporal Decoder and re-format it for a computer or display system. The details of this formatting will vary between applications. In a simple... Token. The DATA Token can have as many bits as are necessary for carrying out processing at a particular place in the system. All other Tokens ignore the extra bits.

A.3.2 The DATA Token

The DATA Token carries data from one processing stage to the next. Consequently, the characteristics of this Token change as it passes through... will be sufficient to collect DATA Tokens and to detect a few Tokens that provide synchronization information (such as PICTURE(underscore)START). In this regard, see subsequent sections A.16, "Connecting... from the data stream. This provides an alternative to doing the configuration via the micro processor interface.

A.3.4 Description of Tokens

This section documents the Tokens which are implemented... 3.5.1. Note: JPEG requires a 2:1:1 structure for its macroblocks when processing 4:2:2 data. See Table A.3.5.

A.3.6 Special Token formats... either is low then the interface is taken to high impedance.

Note: on-chip data processing is not terminated when the DRAM interface is at high impedance. Therefore, errors will occur... decoded video's picture rate. Accordingly, this clock can be used to provide audio/video synchronization.

A.7.1 Spatial Decoder clock signals

The Spatial Decoder has two different (and potentially... in accordance with the present invention, must know what video standard is being input for processing. Thereafter, the system can accept either pre-existing Tokens or raw byte data which is... time a value is written into coded(underscore)data (7:0). Software is responsible for settling

coded(underscore)extn to 0 before the last word of any Token is written to... ..0). The start of this new DATA Token then passes into the Spatial Decoder for processing.

Each time a new 8 bit value is written to coded(underscore)data (7:0... ..Detector analyses data in the DATA Tokens bit serially. The Detector's normal rate of processing is one bit per clock cycle (of coded(underscore)clock). Accordingly, it will typically decode a byte of coded data every 8 cycles of coded(underscore)clock. However, extra processing cycles are occasionally required, e.g., when a non-DATA Token is supplied or when the main decoder(underscore)clock. Data transfer is synchronized to decoder(underscore)clock on-chip.

SECTION A.11 Start code detector

A.11.1... ..Code Detector. So, accessing these registers will be unreliable if the Start Code Detector is processing data. The user is responsible for ensuring that the Start Code Detector is halted before... ..Detector. In this case, the Tokens are passed through the Start Code Detector with no processing to other stages of the Spatial Decoder. These Tokens can only be inserted just before...result will be unpredictable if this is done when the Start Code Detector is actively processing data.

Discard all mode can be safely initiated after any of the Start Code Detector...start code non-alignment interrupt is suppressed.

In contrast, however, JPEG was designed for a computer environment where byte alignment is guaranteed. Therefore, marker codes should only be detected when byte... ..the other hand, was designed to meet the needs of both communications (bit serial) and computer (byte oriented) systems. Start codes in MPEG data should normally be byte aligned. However, the...result will be unpredictable if this is done when the Start Code Detector is actively processing data. So, before initiating a start code search, the Start Code Detector should be stopped so no data is being processed. The Start Code Detector is always in this condition if any of the Start Code...the spatial video decoding circuits (inverse modeler, quantizer and DCT). This second logical buffer allows processing time to include a spread so as to accommodate processing pictures having varying amounts of data.

Both buffers are physically held in a single off 1.1, the unit for all the above mentioned registers is a 512 bit block of data. Accordingly, the... ..until there is space in the buffer. If a buffer continues to be full, more processing stages "up steam" of the buffer will halt until the Spatial Decoder is unable to... ..converting coded data into Tokens started by the Start Code Detector. There are four main processing blocks in the Video Demux: Parser State Machine, Huffman decoder (including an ITOD), Macroblock counter... ..or state machine follows the syntax of the coded video data and instructs the other units. The Huffman decoder converts variable length coded (VLC) data into integers. The Macroblock counter keeps...

Claims: ...said extension indicators, whereby the length of said token can be unlimited ;

an arithmetic logic unit (ALU) ; and

a data buffering means immediately following said system,

whereby time spread for video... ..two-wire interface interconnecting said Huffman decoder with said input shifter, said interface enabling serial processing for data and parallel processing for control; wherein said two-wire interface comprises a sender, a receiver, and a clock... ..and generating command signals, wherein the command signals are communicated to said index to data unit and said arithmetic logic unit for control thereof ;

a ROM accessible to said state machine having separate stored programs for each of a plurality of picture standards, said programs being selectable by a token, whereby processing for a plurality of picture standards is facilitated.

4. The system according to claim 1, further comprising :

a two-wire interface interconnecting processing stages, said interface enabling serial processing for data and parallel processing for control, wherein said two-wire interface comprises: a sender, a receiver, and a clock... ..said programs being selectable by a token ; and

a token formatter for formatting tokens, whereby processing for a plurality of picture standards is facilitated and DATA tokens are created ; wherein said... ..1, further comprising a parser stage for sending an instruction to said Index to Data Unit to select tables needed for a particular identified coding standard, said parser stage indicating whether...

6/K/11 (Item 11 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Start code detecting apparatus for video data stream

Country	Number	Kind	Date
---------	--------	------	------

Abstract ...A2

two successive ones of said processing stages being connected by a two-wire link, wherein said two-wire link comprises: a... ..token being determined by said extension bits; whereby said tokens are unlimited in length ;

said processing stages comprising a spatial decoder accepting an encoded data stream having a plurality of video... ..decoder of said spatial decoder, and responsive to said FLUSH token a portion of said processing stages are reconfigured to await arrival of further data.

Type	Pub. Date	Kind	Text
Available Text	Language	Update	Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: INTRODUCTION

The present invention is directed to improvements in methods and apparatus for decompression which operates to decompress and/or decode a plurality of differently encoded input... ..of the well known standards known as JPEG, MPEG and H.261.

A serial pipeline processing system of the present invention comprises a single two-wire bus used for carrying unique... ..to a plurality of adaptive decompression circuits and the like positioned as a reconfigurable pipeline processor .

PRIOR ART

One prior art system is described in United States Patent No. 5,216,724. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor , dual port memory, scratch-pad memory, and an arbitration mechanism. A first

bus couples the compute modules and a host processor. The device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus.

United States Patent No. 4,785... ...a known quad tree data structure.

United States Patent No. 5,122,875 discloses an apparatus for encoding/decoding an HDTV signal. The apparatus includes a compression circuit responsive to high definition video source signals for providing hierarchically layered... ...compressed video data of relatively greater and lesser importance to image reproduction respectively. A transport processor, responsive to the high and low priority codeword sequences, forms high and low priority transport... ...United States Patent No. 5,168,356 discloses a video signal encoding system that includes apparatus for segmenting encoded video data into transport blocks for signal transmission. ...in respective transport blocks.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening. This is accomplished by an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ...frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ...oscillators and the receiver can continuously receive each channel, then the receiver need not be synchronized with the transmitter. An FFT algorithm implements a fast discrete approximation to the continuous case in which the receiver synchronizes to the first frame and then acquires subsequent frames every frame period. The frame period... ...increasing the amount of data transmitted.

United States Patent No. 5,212,742 discloses an apparatus and method for processing video data for compression/decompression in real-time. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and host processor. Lastly, the device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus. The method handles assigning portions of the image for each of the processors to operate upon.

United States Patent No. 5,231,484 discloses a system and method... ...MPEG standards. Included are three cooperating components or subsystems that operate to variously adaptively pre-process the incoming digital motion video sequences, allocate bits to the pictures in a sequence, and... ...States Patent No. 5,267,334 discloses a method of removing frame redundancy in a computer system for a sequence of moving images. The method comprises detecting a first scene change... ...facing" keyframe or intraframe, and it is normally present in CCITT compressed video data. The process then comprises generating at least one intermediate compressed frame, the at least one intermediate compressed frame containing difference information from the first image for at least one image following...change, known as a "backward-facing" keyframe. The first keyframe and the at least one intermediate compressed frame are linked for forward play, and the second keyframe and the intermediate compressed frames are linked in reverse for reverse play. The intraframe may also be used... ...of complete scene information.

United States Patent No. 5,276,513 discloses a first circuit apparatus, comprising a given number of prior-art image-pyramid stages, together with a second circuit apparatus, comprising the same given number of novel motion-vector stages, perform cost-effective hierarchical motion analysis (HMA) in real-time, with minimum system processing delay and/or employing minimum system processing delay and/or employing minimum hardware structure. Specifically, the first and second circuit apparatus, in response to relatively high-resolution image data from an ongoing input series of successive... ...a relatively high frame rate (e.g., 30 frames per second), derives, after a certain processing-system delay, an ongoing output series of successive given pixel-density vector-data frames that... ...of successive image frames.

United States Patent No. 5,283,646 discloses a method and apparatus for enabling a real-time video encoding system to accurately deliver the desired number of... ..desired bit allocations.

The article, Chong, Yong M., A Data-Flow Architecture for Digital Image Processing, Wescon Technical Papers: No. 2 Oct./Nov. 1984, discloses a real-time signal processing system specifically designed for image processing. More particularly, a token based data-flow architecture is disclosed wherein the tokens are of... ..width having a fixed width address field. The system contains a plurality of identical flow processors connected in a ring fashion. The tokens contain a data field, a control field and a tag. The tag field of the token is further broken down into a processor address field and an identifier field. The processor address field is used to direct the tokens to the correct data-flow processor, and the identifier field is used to label the data such that the data-flow processor knows what to do with the data. In this way, the identifier field acts as an instruction for the data-flow processor. The system directs each token to a specific data-flow processor using a module number (MN). If the MN matches the MN of the particular stage... ..to locate the decoder in the preceding stage in order to pre-decode complex decoding processing and to alleviate critical path problems in the logic circuit. The elastic nature of the...of block signal in most cases.

United States Patent No. 4,903,018 discloses a process and data processing system for compressing and expanding structurally associated multiple data sequences. The process is particular to data sets in which an analysis is made of the structure in... ..data series on the basis of the order number of these data elements. The data processing system for performing the processes includes a storage matrix (26) and an index storage (28) having line addresses of the... ..the final actual video.

United States Patent No. 5,060,242 discloses an image signal processing system DPCM encodes the signal, then Huffman and run length encodes the signal to produce... ..tightly packed without gaps for efficient transmission without loss of any data. The tightly packed apparatus has a barrel shifter with its shift modulus controlled by an accumulator receiving code word... ..OR gate is connected to the shifter, while a register is connected to the gate. Apparatus for processing a tightly packed and decorrelated digital signal has a barrel shifter and accumulator for unpacking... ..an inverse DCPM decoder.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening is accomplished by use of an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ..frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ..kernel matrix.

United States Patent No. 5,231,486 discloses a high definition video system processes a bitstream including high and low priority variable length coded Data words. The coded Data... ..packed High Priority Data and packed Low Priority Data by means of respective data packing units. The coded Data is continuously applied to both packing units. High Priority and Low Priority Length words indicating the bit lengths of high priority and... ..States Patent No. 5,287,178 discloses a video signal encoding system includes a signal processor for segmenting encoded video data into transport blocks having a header section and a packed data section. The system also includes reset control apparatus for releasing resets of system components, after a global system reset, in a prescribed non-simultaneous phased sequence to enable signal processing to commence in the prescribed sequence. The phased reset release sequence begins when valid data...United States Patent No. 5,142,380 to Sakagami et al. discloses an image compression apparatus

suitable for use with still images such as those formed by electronic still cameras using... ..and Q.

United States Patent No. 5,193,002 to Guichard et al. disclosed an apparatus for coding/decoding image signals in real time in conjunction with the CCITT standard H.261. A digital signal processor carries out direct quantization and reverse quantization.

United States Patent No. 5,241,383 to Chen et al. describes an apparatus with a pseudo-constant bit rate video coding achieved by an adjustable quantization parameter. The... relates to an improved pipeline system having an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along the pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the processing stages in the pipeline and interacting with selected stages in the pipeline for control data and/or combined control-data functions among the processing stages, so that the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing. In accordance with the invention, the processing stages may be configurable in response to recognition of at least one token. One of the processing stages may be a Start Code Detector which receives the input and generates and/or... and resetting the system, and a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards. The present invention... data and having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video pictures of varying data size can be controlled. Also in accordance with the invention, a processing stage receives the input data stream, the stage including means for recognizing specified bit stream patterns, whereby the processing stage facilitates random access and error recovery. The invention may also include a means for... invention also includes an inverse modeller stage, an inverse discrete cosine transform stage, and a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for processing data.

In addition, the present invention relates to an improved pipeline system having a Huffman...pipeline stage that incorporates a two-wire transfer control and also shows two consecutive pipeline processing stages with the two-wire transfer control;

Figures. 5a and 5b taken together depict one... shown in Figures. 8a and 8b.

Figure 10 is a block diagram of a reconfigurable processing stage;

Figure 11 is a block diagram of a spatial decoder;

Figure 12 is a... decoder including the prediction filters;

Figure 18 is a pictorial representation of the prediction filtering process;

Figure 19 shows a generalized representation of the macroblock structure;

Figure 20 shows a generalized... buffer;

Figure 25 is a pictorial diagram illustrating prediction data offset from the block being processed;

Figure 26 is a pictorial diagram illustrating prediction data offset by (1,1);

Figure 27...in general terms, the present invention provides an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along a pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the stages in the pipeline and interacting with selected stages in the pipeline for control, data and/or combined control-data functions among the processing stages, whereby the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing.

Each of the processing stages in the pipeline may include both primary and secondary storage, and the stages in processing stages for performance of functions or position independent of the processing stages for performance of functions.

In a pipeline machine, in accordance with the invention, the... ..altered by interfacing with the stages, and the tokens may interact with all of the processing stages in the pipeline or only with some but less than all of said processing stages. The tokens in the pipeline may interact with adjacent processing stages or with non-adjacent processing stages, and the tokens may reconfigure the processing stages. Such tokens may be position dependent for some functions and position independent for other... ..be Huffman coded.

In the improved pipeline machine, the tokens may be generated by a processing stage. Such pipeline tokens may include data for transfer to the processing stages or the tokens may be devoid of data. Some of the tokens may be identified as DATA tokens and provide data to the processing stages in the pipeline, while other tokens are identified as control tokens and only condition the processing stages in the pipeline, such conditioning including reconfiguring of the processing stages. Still other tokens may provide both data and conditioning to the processing stages in the pipeline. Some of said tokens may identify coding standards to the processing stages in the pipeline, whereas other tokens may operate independent of any coding standard among the processing stages. The tokens may be capable of successive alteration by the processing stages in the pipeline.

In accordance with the invention, the interactive flexibility of the tokens in cooperation with the processing stages facilitates greater functional diversity of the processing stages for resident structure in the pipeline, and the flexibility of the tokens facilitates system... ..or alteration. The tokens may be capable of facilitating a plurality of functions within any processing stage in the pipeline. Such pipeline tokens may be either hardware based or software based... ..system bandwidth in the pipeline. The tokens may provide data and control simultaneously to the processing stages in the pipeline.

The invention may include a pipeline processing machine for handling plurality of separately encoded bit streams arranged as a single serial bit... ..and for passing unrecognized control tokens along the pipeline, and a reconfigurable decode and parser processing means responsive to a recognized control token for reconfiguring a particular stage to handle an... ..be a pipeline system and the Start Code Detector may be positioned as the first processing stage in the pipeline.

The present invention also provides, in a system having a plurality of processing stages, a universal adaptation unit in the form of an interactive interfacing token for control and/or data functions among the processing stages, the token being a PICTURE(underscore)START code token for indicating that the start... ..The token may also be a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards.

The CODING(underscore... ..standard as JPEG, and/or any other appropriate picture standard. At least some of the processing stages reconfigure in response to the CODING(underscore)STANDARD token.

One of the processing stages in the system may be a Huffman decoder and parser and, upon receipt of... ..Data stage, and the parser stage may send an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... ..video data, having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video... ..controlled.

The system may include a spatial decoder having a two-wire interface interconnecting processing stages, the interface enabling serial processing for data and parallel processing for control.

As previously indicated, the system may further include a ROM having separate stored... ..of a plurality of picture standards, the programs being selectable by a token to facilitate processing for a plurality of different picture standards.

The spatial decoder system also includes a token... ..decoding stage and a parser stage for sending an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... ..The present invention also provides a pipeline system having an input data stream, and a processing stage for receiving the input data stream, the stage including means for recognizing specified bit... ..whereby said stage facilitates random access and error recovery. In accordance with the invention, the processing stage may be a start code detector and the bit stream patterns may include start... ..token and padding insures uniformity of word size. In accordance with the invention, a reconfigurable processing stage may be provided as a spatial decoder and the padding means adds to picture... ..that if the DATA token has less than the predetermined length, the padder circuit adds units of data to the DATA token until the predetermined length is achieved. A bypass circuit...1 tokens into a buffer, having a second predetermined width.

The invention also provides an apparatus for providing a time delay to a group of compressed pictures, the pictures corresponding to... ..and capable of delaying the words of data, is in communication with a control circuit intermediate the counter circuit and the inverse modeller circuit, the control circuit also communicating with the... ..inverse modeller stage and an inverse discrete cosine transform stage, the improvement characterized by a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for

processing data.

In accordance with the invention, the token may be a QUANT(underscore)TABLE token for causing the processing stage to generate a quantization table.

The present invention also provides a Huffman decoder for... ..of bits used to represent an item of data.

DECODER: An embodiment of a decoding process.

DECODING (PROCESS): The process defined in this specification that reads an input coded bitstream and produces decoded pictures or... ..the same order in which they were presented at the input of the encoder.

ENCODING (PROCESS): A process, not specified in this specification, that reads a stream of input pictures or audio samples...to provide an estimate of the pel value or data element currently being decoded.

RECONFIGURABLE PROCESS STAGE (RPS): A stage, which in response to a recognized token, reconfigures itself to perform various operations.

SLICE: A series of macroblocks.

TOKEN: A universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions... ..indicates that the corresponding stage holds valid data, i.e., data that is to be processed in one of the pipeline stages. After processing (which may involve nothing more than a simple transfer without manipulation of the data) valid... ..present invention may be used with any number of pipeline stages. Furthermore, data may be processed in more than one stage and the processing time for different stages can differ.

In addition to clock and data signals (described below... ..other system. For example, the last pipeline stage may pass its data on to subsequent processing circuitry. The ACCEPT signal, which is illustrated as the lower of the two lines connecting... ..the minimum disturbance possible to other pipeline stages. Succeeding pipeline stages are allowed to continue processing and, therefore, this means that gaps open up in the stream of data following the... ..The data in the pipeline is encoded such that many different types of data are processed in the pipeline. This encoding accommodates data packets of variable size and the size of...the other hand, it may generate itself, all or part of the data to be processed in the pipeline. Indeed, as is explained below, a "stage" may contain arbitrary

processing circuitry, including none at all (for simple passing of data) or entire systems (for example... values zero and 255 may not be used.

If such a picture were to be processed in a pipeline built in the practice of the present invention, then one of these... data must not be written over since it is data that must be saved for processing or use in a downstream device e.g., a pipeline stage, a device or a... connected to the pipeline upstream contains data D4 that is to be transferred into and processed in the pipeline. ...pipeline, in accordance with the preferred embodiments of the present invention, to "fill up" empty processing stages is highly advantageous since the processing stages in the pipeline thereby become decouple from one another. In other words, even though... data can be transferred into the pipeline and between stages even when one or more processing stages is blocked.

In the embodiment shown in Fig. 1, it is assumed that the... propagate all the way back to the beginning of the pipeline if there is some intermediate stage that is able to accept new data.

In the embodiment illustrated in Fig. 1... has been mentioned. It is to be further understood that each pipeline stage may also process the data it has received arbitrarily before passing it between its internal storage elements or... the portion of the pipeline that contains input and output storage elements and that arbitrarily processes data stored in its storage elements.

Furthermore, the "device" ... valid data, but also when a stage requires more than one clock phase to finish processing its data. This also can occur when it creates valid data in one or both... control the passage of data between adjacent storage elements. The VALID signal may also be processed in an analogous manner.

A great advantage of the two-wire interface (one wire for... In addition, two extra latches and a small number of gates are preferably added to process the ACCEPT and VALID signals that are associated with the data latches in each half... application so requires. The interface in accordance with this embodiment can also be used to process analog signals.

As discussed previously, while other conventional timing arrangements may be used, the interface... circuit B1, which may be provided to convert output data from input latch LDIN into intermediate data, which is then later loaded in an output data latch LDOUT, which comprises the... is connected either directly as an input to the validation output latch LVOUT, or via intermediate logic devices or circuits that may alter the signal.

Similarly, the output validation signal QVOUT... to the input of the validation input latch QVIN of the following stage, or via intermediate devices or logic circuits, which may alter the validation signal. This ...word.

Preferred Data Structure - "tokens"

In the sample application shown in Fig. 4, each stage processes all input data, since there is no control circuitry that excludes any stage from allowing... are connected together in a relatively simple configuration. The simplest configuration is a pipeline of processing steps. For example, in the one shown in Fig. 1. The use of tokens, however... flows from left to right in the diagram. Data enters the machine and passes into processing Stage A. This may or may not modify the data and it then passes the... advantage of the tokens is their ability to achieve this kind of communication. Since any processing stage that does not recognize a token simply passes it on unaltered to the next... is transmitted along with the address and data fields in each token so that a processing stage can pass on a token (which can be of arbitrary length) without having to... be the first word of a new token.

Note that although the simple pipeline of processing stages is particularly useful, it will be appreciated that tokens may be applied to more complicated configurations of processing elements. An example of a more complicated processing element is described below.

It is not necessary, in accordance with the present invention, to... has extension bits. An example of this is a token that activates a stage that processes video quantization values stored in a quantization table (typically a memory

device). For example, a...turn, is of great importance in video data pipeline systems since it ensures that all processing stages can be continuously running at full bandwidth.

In accordance to the present invention, in... some other chips in the set. This is advantageous both from the perspective of a customer and from that of a chip manufacturer. Even if modifications mean that all chips are...the end of a token (and hence the start of the next token) to be processed correctly (including simple non-manipulative transfer), even if the token is not recognized by the... block diagram of a pipeline stage whose function is as follows. If the stage is processing a predetermined token (known in this example as the DATA token), then it will duplicate... the address field of the DATA token. If, on the other hand, the stage is processing any other kind of token, it will delete every word. The overall effect is that... respective output signals:

In the duplication stage, the output from the data latch LDIN forms intermediate data referred to as MID(underscore)DATA. This intermediate data word is loaded into the data output latch LDOUT only when an intermediate acceptance signal (labeled "MID(underscore)ACCEPT" in Fig. 8a) is set HIGH.

The portion of... data. These include a "DATA(underscore)TOKEN" signal that indicates that the circuitry is currently processing a valid DATA Token, and a NOT(underscore)DUPLICATE signal which is used to control duplication of data. When the circuitry is processing a DATA Token, the NOT(underscore)DUPLICATE signal toggles between a HIGH and a LOW... the token to be duplicated once (but no more times). When the circuitry is not processing a valid DATA Token then the NOT(underscore)DUPLICATE signal is held in a HIGH state. Accordingly, this means that the token words that are being processed are not duplicated.

As Fig. 8a illustrates, the upper six bits of 8-bit intermediate data word and the output signal QI1 from the latch LI1 form inputs to a... explained further below.

Latch LO1 performs the function of latching the last value of the intermediate extension bit (labeled "MID(underscore)EXTN" and as signal S4), and it loads this value... and the DATA(underscore)TOKEN signal will become "0", indicating that the circuitry is not processing a DATA token.

If QI1 is "0" and SO is "0", thereby indicating a DATA... phase and the DATA(underscore)TOKEN signal will become "1", indicating that the circuitry is processing a DATA token.

The NOT(underscore)DUPLICATE signal (the output signal QO3) is similarly loaded...LVOUT at the same time that MID(underscore)DATA is loaded into LDOUT and the intermediate extension bit (signal S4) is loaded into LEOUT. Signal S5 is also combined with the...above. This has the effect that all tokens except the one that causes the duplication process will be deleted from the token stream, since a device connected to the output terminals (OUTDATA, OUTEXTN and OUTVALID) will not recognize these token words as valid data.

As before... and is duplicated.

Referring now more particularly to Figure 10, there is shown a reconfigurable process stage in accordance with one aspect of the present invention.

Input latches 34 receive an... the input latches 34 is passed as a first input over line 35 to a processing unit 36. A first output from the token decode subsystem 33 is passed over line 37 as a second input to the

processing unit 36. A second output from the token decode 33 is passed over line 40 to an action identification unit 39. The action identification unit 39 also receives input from registers 43 and 44 over line 46. The registers 43... is determined by the history of tokens previously received. The output from the action identification unit 39 is passed over line 38 as a third input to the processing unit 36. The output from the processing unit 36 is passed to output latches 41. The output from the output latches 41 is... decoder 56 is passed over line 63 as an input to an Index to Data Unit (ITOD) 64. The Huffman decoder 56 and the ITOD 64 work together as a single logical unit.

The output from the ITOD 64 is passed over line 65 to an arithmetic logic unit (ALU) 66. A first output from the ALU 66 is passed over line 67 to...blocks 133.

Referring to Figure 14b, in the JPEG and H.261 standards, the Common Intermediate Format (CIF) is used, wherein a picture 141 is encoded as 6 rows each containing... ..in a zigzag direction indicated by the arrow 144. The GOBs 142 are, in turn, processed row-by-row, left-to-right in each row.

Referring now to Figure 14c, it... ..in accordance with the practice of the present invention. A first picture 161 to be processed contains a first PICTURE(underscore)START token 162, first-picture information of indeterminate length 163, and a first PICTURE(underscore)END token 164. A second picture 165 to be processed contains a second PICTURE(underscore)START token 166, second picture information of indeterminate length 167... ..tokens 162 and 166 indicate the start of the pictures 161 and 165 to the processor. Likewise, the PICTURE(underscore)END tokens 164 and 168 signify the end of the pictures 161 and 165 to the processor. This allows the processor to process picture information 163 and 167 of variable lengths.

Referring to Figure 17, a split 171...Video Formatter (not shown in Figure 17).

Referring now to Figure 18, the prediction filtering process is illustrated. A forward picture 201 is passed over line 202 as a first input... ..the right of the value decode shift register 230, as indicated by area 231. This process eliminates overlapping start code images, as discussed below. A first output from the value decode... ..Code Detector. The Start Code Detector then receives a first data value image 244. Before processing the first data value image 244, the Start Code Detector may detect a second start... ..image 244 at a length 246. If this occurs, the Start Code Detector does not process the first data value image 244, and instead receives and processes a second data value image 247.

...line 1 of Table 600, whenever a "sequence start" image is received during H.261 processing or a "picture start" image is received during MPEG processing, the entire group of four control tokens is generated, each followed by its corresponding data...Picture Decoding

3. Motion Picture Decompression

4. RAM Memory Map

5. Bitstream Characteristics

6. Reconfigurable Processing Stage

7. Multi-Standard Coding

8. Multi-Standard Processing Circuit-2nd Mode of Operation

9. Start Code Detector

10. Tokens

11. DRAM Interface

12... ..described herein in greater detail) and reformatting this output for use, including display in a computer or other display systems, including a video display system. Implementation of this formatting varies significantly... ..the Spatial Decoder circuits.

The Spatial Decoder of the present invention performs all the required processing within a single picture. This reduces the redundancy within one picture.

The Temporal Decoder reduces... ..modeller 75, the inverse zig-zag 81 and the inverse DCT 83. The standard independent units within the Huffman decoder and parser include the ALU 66 and the token formatter 71.

Referring now to Figure 12, the standard-independent units include the DRAM interface 100, the fork 91, the FIFO register 96, the summer 98 and the output selector 106. The standard dependent units are the address generator 94, which is different in H.261 and in MPEG, and...much of the operation is very similar between the three different compression standards.

The next unit is the state machine 68 (Figure 11) located within the Huffman decoder and parser. Here... The same holds true for JPEG, which is a third, completely independent program.

The next unit is the Huffman decoder 56 which functions with the index to data unit 64. Those two units cooperate together to perform the Huffman decoding. Here, the algorithm that is used for Huffman... to the Huffman decoder at different times consistent with the standard in operation.

The last unit on the chip that is dependent on the compression standard is the inverse quantizer 79...an H.261 group of blocks and an MPEG slice. When H.261 data is processed after the Start Code Detector, each group of blocks is preceded by a slice(underscore... these standards have totally different sets of tables.

As previously indicated, most of the system units are compression standard independent. If a unit is standard independent, and such units need not remember what CODING(underscore)STANDARD is being processed. All of the units that are standard dependent remember the compression standard as the CODING(underscore)STANDARD token flows... CODING(underscore)STANDARD tokens at the Start Code Detector that is positioned as the first unit in the pipeline, this change of compression standard is readily handled. The token says a... found in the standard, i.e. from the bitstream into a prediction mode token. This processing is performed by the Huffman decoder and parser state machine, where it is easy to... to that token. By having these tokens and using them appropriately, the design of other units in the machine is simplified. Although there may be some complications in the program, benefits... a first encoded signal (the MPEG or H.261 encoded video signal) in a pipeline processing system. The Temporal Decoder is not needed for JPEG decoding.

In this regard, the invention... the use of a single pipeline decoder and decompression system. The decoding and decompression pipeline processor is organized on a unique and special configuration which allows the handling of the multi... video signals through the use of techniques all compatible with the single pipeline decoder and processing system. The Spatial Decoder is combined with the Temporal Decoder, and the Video Formatter is...with only still pictures. The compression standard independent Spatial Decoder performs all of the data processing within the boundaries of a single picture. Such a decoder handles the spatial decompression of... to the multi-standard, configurable Video Formatter, which then provides an output to the display terminal. In a first sequence of similar pictures, each decompressed picture at the output of the... of control tokens and DATA tokens, in combination with a plurality of sequentially-positioned reconfigurable processing stages selected and organized to act as a standard-independent, reconfigurable-pipeline-processor.

With regard to JPEG decoding, a single Spatial Decoder with no off chip DRAM can... video. Accordingly, signals carried by DATA tokens pass directly through the Temporal Decoder without further processing when the Temporal Decoder is configured for a JPEG operation.

Another aspect of the present... for subsequent use in temporal decoding of subsequent pictures.

Generally, the Temporal Decoder performs the processing between pictures either earlier and/or later in time with reference to the picture currently... is distributed among several areas of DRAM in the sense that the decompressed output information, processed by the Spatial Decoder, is stored in other DRAM registers by other random access memories...first decoder circuit (the Spatial Decoder) directly to the Video Formatter for handling without signal processing delay.

The Temporal Decoder also reorders the blocks of picture data for display by a... from a selection of pictures which have arrived earlier or later than the picture under processing. When a picture is described in this context, it may

mean any one of the... 2. The result, i.e., the final decoded picture resulting from the addition of a process step performed by the decoder;

3. Previously decoded pictures read from the DRAM; and

4... START token and a subsequent PICTURE(underscore)END token.

After the picture data information is processed by the Temporal Decoder, it is either displayed or written back into a picture memory location. This information is then kept for further reference to be used in processing another different coded data picture.

Re-ordering of the MPEG encoded pictures for visual display...used to encode a referenced picture of a picture might be identified as being one unit long, another picture might be a number of units long, while still a third picture could be a fraction of that unit.

None of the existing standards (MPEG 1.2, JPEG, H.261) define a way of... picture rate, whereas the Video Formatter can handle a variable input picture rate.

6. RECONFIGURABLE PROCESSING STAGE

Referring again to Figure 10, the reconfigurable processing stage (RPS) comprises a token decode circuit 33 which is employed to receive the tokens... input latches 34. The output of the token decode circuit 33 is applied to a processing unit 36 over the two-wire interface 37 and an action identification circuit 39. The processing unit 36 is suitable for processing data under the control of the action identification circuit 39. After the processing is completed, the processing unit 36 connects such completed signals to the output, two-wire interface bus 40 through output... token decode circuit 33 are applied simultaneously to the action identification circuit 39 and the processing unit 36. The action identification function as well as the RPS is described in further detail... not standard independent circuits. The data flows through the token decode circuit 33, through the

processing unit 36 and onto the two-wire interface circuit 42 through the output latches 41. If... wire interface 42 through the output circuit 41. The present invention operates as a pipeline processor having a two-wire interface for controlling the movement of control tokens through the pipeline... time, the token decode circuit 33 provides a proper flag or index signal to the processing unit 36 to alert it to the presence of the token being handled by the action identification circuit 39.

Control tokens may also be processed.

A more detailed description of the various types of tokens usable in the present invention... standard now passing through the state machine shown with reference to Figure 10.

Similarly, the processing unit 36 which is under the control of the action identification circuit 39 is now ready to process the information contained in the data fields of the DATA token when it is appropriate... action identification circuit 39 and is immediately followed by a DATA token which is then processed by the processing unit 36. The control token exits the output latches circuit 41 over the output two-wire interface 42 immediately preceding the DATA token which has been processed within the processing unit 36.

In the present invention, the action identification circuit, 39, is a state machine holding... show that the action can also be affected by the token that is currently being processed by the token decode circuit 33.

In general, there is shown token decoding and data processing in accordance with the present invention. The data processing is performed as configured by the action identification circuit 39. The action is affected by... information stored from previously decoded tokens in registers 43 and 44, the current token under processing, and

the state and history information that the action identification unit 39 has itself acquired. A distinction is thereby shown between Control tokens and DATA tokens.

In any RPS, some tokens are viewed by that RPS unit as being Control tokens in that they affect the operation of the RPS presumably at... are viewed by the RPS as DATA tokens. Such DATA tokens contain information which is processed by the RPS in a way that is determined by the design of the particular... view of the same token. Some of the tokens might be viewed by one RPS unit as DATA Tokens while another RPS unit might decide that it is actually a Control Token. For example, the quantization table information... into a token called a quantization table token (QUANT(TABLE)) which goes down the processing pipeline. As far as that machine is concerned, all of that was data; it was... sort of data into another sort of data, which is clearly a function of the processing performed by that portion of the machine. However, when that information gets to the inverse... present. This information is viewed as control information, and then that control information affects the processing that is done on subsequent DATA tokens because it affects the number that you multiply... important feature of the invention is that each of the stages of circuitry has the processing capability within it to be able to perform the necessary operations for each of the... operations are to be performed at a given time, come as tokens. There is one processing element that differs between the different stages to provide this capability. In the state machine... standard is and it looks up the parameters that it needs to apply to the processing elements in order to perform a proper operation. For example, the inverse quantizer will look... is set to 1 for a particular compression standard, and will apply that to its processing circuitry.

In a similar sense the Huffman decoder 56 has a number of tables within... MPEG video standard or the JPEG video standard. These three compression coding standards specify similar processes to be done on the arriving data, but the structure of the datastreams is different... token stream embodying the current coding standard. The control tokens are passed through the pipeline processor, and are used, i.e., decoded, in the state machines to which they are relevant... this regard, the DATA Tokens are treated in the same fashion, insofar as they are processed only in the state machines that are configurable by the control tokens into processing such DATA Tokens. In the remaining state machines, they pass through unchanged.

More specifically, a... signals. The remaining portions of the token are used to indicate and identify the internal processing control function which is standard for all of the datastreams passing through the pipeline processor. In one form of the invention, the token extension is used to carry the current... accompanying data. As previously discussed, this information is utilized in the system to reconfigure the processing stage used to perform the function required by the various standards created for that purpose... picture number as indicated by the value.

The system also includes a multi-stage parallel processing pipeline operating under the principles of the two-wire interface previously described. Each of the... the token presently entering the state machine into the action identification circuit 39 or the processing unit 36, as appropriate. The processing unit has been previously reconfigured by the next previous control token into the form needed for handling the current coding standard, which is now entering the processing stage and carried by the next DATA token. Further, in accordance with this aspect of the invention, the succeeding state machines in the processing pipeline can be functioning under one coding standard, i.e., H.261, while a previous... tokens required to decode a number of coding standards with a fixed number of reconfigurable processing stages. More specifically, the PICTURE(END) control token is employed because it is important... standard machine, it is necessary to create additional control tokens within the multi-standard pipeline processing machine which will then indicate which one of the standard decoding techniques to use. Such... and to push the current picture through the decoder to the display.

8. MULTI-STANDARD PROCESSING CIRCUIT - SECOND MODE OF OPERATION

A compression standard-dependent circuit, in the form of the... of the Start Code Detector will subsequently be discussed in further detail, as will the process of starting up of the decoder.

The aforementioned description has been concerned primarily ...the data which immediately follows according to the standard. However, in the multi-standard pipeline processing system of the present invention, where compatibility is required for multiple standards, the system has... ..signals, including flag signals, are generated by each state machine to handle some of the processing within that state machine. Values carried in the standards can be used to access machine... ..its contents must be removed from the two wire interface to ensure that no further processing takes place using these 3 bytes. The decode register is emptied, and the value decode... ..10. TOKENS

In the practice of the present invention, a token is a universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions and is adapted for use with a reconfigurable processing stage (RPS) which is a stage, which in response to a recognized token, reconfigures itself to perform various operations.

Tokens may be either position dependent or position independent upon the processing stages for performance of various functions. Tokens may also be metamorphic in that they can be altered by a processing stage and then passed down the pipeline for performance of further functions. Tokens may interact... ..other functions, and the specific interaction with a stage may be conditioned by the previous processing history of a stage.

A PICTURE(underscore)END token is a way of signalling the... ..through a fixed size, fixed width buffer.

The present invention is directed to a pipeline processing system which has a variable configuration which uses tokens and a two-wire system. The... ..do not use control tokens.

The control tokens are generated by circuitry within the decoder processor and emulate the operation of a number of different type standard-dependent signals passing into the serial pipeline processor for handling. The technique used is to study all the parameters of the multi-standards that are selected for processing by the serial processor and noting 1) their similarities, 2) their dissimilarities, 3) their needs and requirements and 4) selecting the correct token function to effectively process all of the standard signals sent into the serial processor. The functions of the tokens are to emulate the standards. A control token function is... ..the standard dependent signals and as an element to transmit control information through the pipeline processor.

In prior art system, a dedicated machine is designed according to well-known techniques to... ..tokens provide and make a sensible format for communicating information through the decompression circuit pipeline processor. In the design selected hereinafter and used in the preferred embodiment, each word of a... ..However, this is not a limitation on the invention, but on the magnitude of the processing steps elected to be accomplished by use of these tokens. It is to be noted... ..bit address for use in accessing the random access memories used throughout this serial decompression processor. This provides an additional degree of variability that facilitates a broad range of versatility.

As previously described, the DATA token carries data from one processing stage to the next. Consequently, the characteristics of this token change as it passes through... ..longest number of data bits because it needs to provide the most information to the

processing unit so that it can start the decompression with as much information as possible. Words which...to receive an address, it waits for the address generator to supply a valid address, processes that address and then sets the accept line high for one clock period. Thus, it... ..be read. This signal passes between two asynchronous clock regimes and, therefore, passes through three synchronizing flip flops.

Provided RAM2 312 is empty, the next item of data to arrive on...interesting.

In general, prediction data will be offset from the position of the block being processed as specified in the motion vectors in x and y. Thus, the block of data... ..address, 9. Data is read from this address and the x value is incremented. The process is repeated until the x value reaches its stop value, at which point, the y... ..is read, the x value is again incremented until it reaches its stop value. The process is repeated until both x and y values have

reached their stop values. Thus, the...invention, is that additional information must be provided to the prediction filters to indicate what processing is required on the data. This consists of the following:

a "last byte" signal indicating... ..bit 0) is incremented and the x address (3 LSBS) is reset to zero. This process is repeated until 64 bytes have been read. With a 16 or 32 bit wide...register while its access register is set to zero, the results are undefined.

14. MICRO-PROCESSOR INTERFACE

A standard byte wide micro-processor interface (MPI) is used on all circuits with in the Spatial Decoder and Temporal Decoder... ..the parameter column. The actual specifications are shown in the respective columns min, max and units.

The DC operating conditions can be seen with reference to Table A.6.3. Here... ..the signal is present the maximum amount of time that this signal is available. The Units column gives the units of measurement used to describe the signals.

16. MPI WRITE TIMING

The general description of...a PICTURE(underscore)END token is decoded and forces the data in the coded data buffers to be applied to the Huffman decoder and video demultiplexor, the final picture can be... ..Consequently, the machine will not go into error recovery mode and will successfully continue to process the coded data.

A still further advantage of the use of a PICTURE(underscore)END token is that the serial pipeline processor will continue the processing of uninterrupted data. Through the use of a PICTURE(underscore)END token, the serial pipeline processor is configured to handle less than the expected amount of data and, therefore, continues processing. Typically, a prior art machine would stop itself because of an error condition. As previously... ..of the Huffman decode and Video Demultiplexor know the number of blocks that it will process during each picture recovery cycle. When the correct number of blocks do not arrive from... ..Each of the state machines recognizes a FLUSH control token as information not to be processed. Accordingly, the FLUSH token is used to fill up all of the remaining empty parts... ..Huffman Decoder and Video Demultiplexor. In this way, the FLUSH token is like padding for buffers.

The Token Decoder in the Huffman circuit recognizes the FLUSH token and ignores the pseudo... ..less information than normally expected to decode the last picture. The Huffman decode circuit finishes processing the information contained in the last picture, and outputs this information through the DRAM interface... ..token, in accordance with the present invention, is used to pass through the entire pipeline processor and to ensure that the buffers are emptied and that other circuits are reconfigured to... ..underscore)END token, a padding word and a FLUSH token indicating to the serial pipeline processor that the picture processing for the current picture form is completed. Thereafter, the various state machines need reconfiguring to... ..FLUSH token resets each stage as it passes through, but-allows subsequent stages to continue processing. This prevents a loss of data. In other words, the FLUSH token is a variable... ..AFTER PICTURE

The STOP(underscore)AFTER(underscore)PICTURE function is employed to shut down the processing of the serial pipeline decompressing circuit at a logical point in its operation. At this... ..a picture, the STOP(underscore)AFTER(underscore)PICTURE operation signals the end of all current processing.

22. MULTI(underscore)STANDARD - SEARCH MODE

Another feature of the present invention is the use... ..underscore)MODE control token which is used to reconfigure the input to the serial pipeline processor to look at the incoming bit stream. When the search mode is set, the Start... ..combination of control tokens, and DATA tokens along with the reconfiguration circuits, to provide similar processing.

The use of search mode in the present invention is convenient in many situations including... video disc. In general, a search mode is convenient when the user interrupts the normal processing of the serial pipeline at a point where the machine does not expect such an... be the case.

In brief, the Huffman Decoder 321 works in conjunction with the other units shown in Figure 27. These other units are the Parser State Machine 322, the inshifter 323, the Index to Data unit 324, the ALU 325, and the Token Formatter 326. As described previously, connection between these blocks is governed by a two wire interface. A more detailed description of how these units function is subsequently described herein in greater detail, the focus here is on particular aspects... control certain functions of the Index to Data 324 and ALU 325. Control of these units by the Huffman Decoder is necessary for proper decoding of block-level information. Having the... further detail in the "More Detailed Description of the Invention" section.

The Index to Data unit 324 performs the second part of the multi-part algorithm. This unit contains a look up table that provides the actual Huffman decoded data. Entries in the... by detecting these in the Huffman Decoder 321, rather than in the Index to Data unit 324.

This index number is then passed to the Index to Data unit 324. In essence, the Index to Data unit is a look-up table. In accordance with one aspect of the algorithm, the look... format that JPEG specifies for transferring an alternate JPEG table.

From the Index to Data unit 324, the decoded index number or other data is passed, together with the accompanying control... the entering data to ensure that the DATA tokens are of the correct size for processing. In fact, the token stream can be corrected in some situations if the error is... an order that is useful for the decompression circuits, but not for the particular display unit being used. When a block of data enters the Buffer Manager, the Buffer Manager supplies... the output of the Spatial Decoder or Temporal Decoder and re-format it for a computer or display system. The details of this formatting will vary between applications. In a simple... Token. The DATA Token can have as many bits as are necessary for carrying out processing at a particular place in the system. All other Tokens ignore the extra bits.

A.3.2 The DATA Token

The DATA Token carries data from one processing stage to the next. Consequently, the characteristics of this Token change as it passes through... will be sufficient to collect DATA Tokens and to detect a few Tokens that provide synchronization information (such as PICTURE(underscore)START). In this regard, see subsequent sections A.16, "Connecting... from the data stream. This provides an alternative to doing the configuration via the micro processor interface.

A.3.4 Description of Tokens

This section documents the Tokens which are implemented... 3.5.1. Note: JPEG requires a 2:1:1 structure for its macroblocks when processing 4:2:2 data. See Table A.3.5.

A.3.6 Special Token formats... either is low then the interface is taken to high impedance.

Note: on-chip data processing is not terminated when the DRAM interface is at high impedance. Therefore, errors will occur... decoded video's picture rate. Accordingly, this clock can be used to provide audio/video synchronization.

A.7.1 Spatial Decoder clock signals

The Spatial Decoder has two different (and potentially... in accordance with the present invention, must know what video standard is being input for processing. Thereafter, the system can accept either pre-existing Tokens or raw byte data which is... time a value is written into coded(underscore)data (7:0). Software is responsible for settling

coded(underscore)extn to 0 before the last word of any Token is written to... ..0). The start of this new DATA Token then passes into the Spatial Decoder for processing.

Each time a new 8 bit value is written to coded(underscore)data (7:0... ..Detector analyses data in the DATA Tokens bit serially. The Detector's normal rate of processing is one bit per clock cycle (of coded(underscore)clock). Accordingly, it will typically decode a byte of coded data every 8 cycles of coded(underscore)clock. However, extra processing cycles are occasionally required, e.g., when a non-DATA Token is supplied or when the main decoder(underscore)clock. Data transfer is synchronized to decoder(underscore)clock on-chip.

SECTION A.11 Start code detector

A.11.1... ..Code Detector. So, accessing these registers will be unreliable if the Start Code Detector is processing data. The user is responsible for ensuring that the Start Code Detector is halted before... ..Detector. In this case, the Tokens are passed through the Start Code Detector with no processing to other stages of the Spatial Decoder. These Tokens can only be inserted just before...result will be unpredictable if this is done when the Start Code Detector is actively processing data.

Discard all mode can be safely initiated after any of the Start Code Detector...start code non-alignment interrupt is suppressed.

In contrast, however, JPEG was designed for a computer environment where byte alignment is guaranteed. Therefore, marker codes should only be detected when byte... ..the other hand, was designed to meet the needs of both communications (bit serial) and computer (byte oriented) systems. Start codes in MPEG data should normally be byte aligned. However, the...result will be unpredictable if this is done when the Start Code Detector is actively processing data. So, before initiating a start code search, the Start Code Detector should be stopped so no data is being processed. The Start Code Detector is always in this condition if any of the Start Code...the spatial video decoding circuits (inverse modeler, quantizer and DCT). This second logical buffer allows processing time to include a spread so as to accommodate processing pictures having varying amounts of data.

Both buffers are physically held in a single off 1.1, the unit for all the above mentioned registers is a 512 bit block of data. Accordingly, the... ..until there is space in the buffer. If a buffer continues to be full, more processing stages "up steam" of the buffer will halt until the Spatial Decoder is unable to... ..converting coded data into Tokens started by the Start Code Detector. There are four main processing blocks in the Video Demux: Parser State Machine, Huffman decoder (including an ITOD), Macroblock counter... ..or state machine follows the syntax of the coded video data and instructs the other units. The Huffman decoder converts variable length coded (VLC) data into integers. The Macroblock counter keeps...

Claims:

1. A pipeline machine, comprising a plurality of processing stages, characterized by:

two successive ones of said processing stages being connected by a two-wire link, wherein said two-wire link comprises: a... ..token being determined by said extension bits; whereby said tokens are unlimited in length;

said processing stages comprising a spatial decoder accepting an encoded data stream having a plurality of video... ..decoder of said spatial decoder, and responsive to said FLUSH token a portion of said processing stages are reconfigured to await arrival of further data.

2. The machine according to claim 1, wherein at least one of said processing stages has a variable length DATA token stored therein, and responsive to said PICTURE(underscore)END token, said one processing stage adds bits to a last word of said DATA token until said DATA token...

6/K/12 (Item 12 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Country	Number	Kind	Date
---------	--------	------	------

Abstract ...A1

Type	Pub. Date	Kind	Text
Available Text	Language	Update	Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: INTRODUCTION

The present invention is directed to improvements in methods and apparatus for decompression which operates to decompress and/or decode a plurality of differently encoded input... ..of the well known standards known as JPEG, MPEG and H.261.

A serial pipeline processing system of the present invention comprises a single two-wire bus used for carrying unique... ..to a plurality of adaptive decompression circuits and the like positioned as a reconfigurable pipeline processor.

PRIOR ART

One prior art system is described in United States Patent No. 5,216,724. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and a host processor. The device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus.

United States Patent No. 4,785... ..a known quad tree data structure.

United States Patent No. 5,122,875 discloses an apparatus for encoding/decoding an HDTV signal. The apparatus includes a compression circuit responsive to high definition video source signals for providing hierarchically layered... ..compressed video data of relatively greater and lesser importance to image reproduction respectively. A transport processor, responsive to the high and low priority codeword sequences, forms high and low priority transport... ..United States Patent No. 5,168,356 discloses a video signal encoding system that includes apparatus for segmenting encoded video data into transport blocks for signal transmission. ...in respective transport blocks.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening. This is accomplished by an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ..frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by

a savings of designated samples, and a... ..oscillators and the receiver can continuously receive each channel, then the receiver need not be synchronized with the transmitter. An FFT algorithm implements a fast discrete approximation to the continuous case in which the receiver synchronizes to the first frame and then acquires subsequent frames every frame period. The frame period... ..increasing the amount of data transmitted.

United States Patent No. 5,212,742 discloses an apparatus and method for processing video data for compression/decompression in real-time. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and host processor. Lastly, the device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus. The method handles assigning portions of the image for each of the processors to operate upon.

United States Patent No. 5,231,484 discloses a system and method... ..MPEG standards. Included are three cooperating components or subsystems that operate to variously adaptively pre-process the incoming digital motion video sequences, allocate bits to the pictures in a sequence, and... ..States Patent No. 5,267,334 discloses a method of removing frame redundancy in a computer system for a sequence of moving images. The method comprises detecting a first scene change... ..facing" keyframe or intraframe, and it is normally present in CCITT compressed video data. The process then comprises generating at least one intermediate compressed frame, the at least one intermediate compressed frame containing difference information from the first image for at least one image following...change, known as a "backward-facing" keyframe. The first keyframe and the at least one intermediate compressed frame are linked for forward play, and the second keyframe and the intermediate compressed frames are linked in reverse for reverse play. The intraframe may also be used... ..of complete scene information.

United States Patent No. 5,276,513 discloses a first circuit apparatus, comprising a given number of prior-art image-pyramid stages, together with a second circuit apparatus, comprising the same given number of novel motion-vector stages, perform cost-effective hierarchical motion analysis (HMA) in real-time, with minimum system processing delay and/or employing minimum system processing delay and/or employing minimum hardware structure. Specifically, the first and second circuit apparatus, in response to relatively high-resolution image data from an ongoing input series of successive... ..a relatively high frame rate (e.g., 30 frames per second), derives, after a certain processing-system delay, an ongoing output series of successive given pixel-density vector-data frames that... ..of successive image frames.

United States Patent No. 5,283,646 discloses a method and apparatus for enabling a real-time video encoding system to accurately deliver the desired number of... ..desired bit allocations.

The article, Chong, Yong M., A Data-Flow Architecture for Digital Image Processing, Wescon Technical Papers: No. 2 Oct./Nov. 1984, discloses a real-time signal processing system specifically designed for image processing. More particularly, a token based data-flow architecture is disclosed wherein the tokens are of... ..width having a fixed width address field. The system contains a plurality of identical flow processors connected in a ring fashion. The tokens contain a data field, a control field and a tag. The tag field of the token is further broken down into a processor address field and an identifier field. The processor address field is used to direct the tokens to the correct data-flow processor, and the identifier field is used to label the data such that the data-flow processor knows what to do with the data. In this way, the identifier field acts as an instruction for the data-flow processor. The system directs each token to a specific data-flow processor using a module number (MN). If the MN matches the MN of the particular stage... ..to locate the decoder in the preceding stage in order to pre-decode complex decoding processing and to alleviate critical path problems in the logic circuit. The elastic nature of the...of block signal in most cases.

United States Patent No. 4,903,018 discloses a process and data processing system for compressing and expanding structurally associated multiple data sequences. The process is particular to data sets in which an analysis is made of the structure in order... ..data series on the basis of the order number of these data elements. The data processing

system for performing the processes includes a storage matrix (26) and an index storage (28) having line addresses of the... ..the final actual video.

United States Patent No.5,060,242 discloses an image signal processing system DPCM encodes the signal, then Huffman and run length encodes the signal to produce... ..tightly packed without gaps for efficient transmission without loss of any data. The tightly packed apparatus has a barrel shifter with its shift modulus controlled by an accumulator receiving code word... ..OR gate is connected to the shifter, while a register is connected to the gate. Apparatus for processing a tightly packed and decorrelated digital signal has a barrel shifter and accumulator for unpacking... ..an inverse DPCM decoder.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of declination, interpolation, and sharpening is accomplished by use of an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ..frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ..kernel matrix.

United States Patent No. 5,231,486 discloses a high definition video system processes a bitstream including high and low priority variable length coded Data words. The coded Data... ..packed High Priority Data and packed Low Priority Data by means of respective data packing units. The coded Data is continuously applied to both packing units. High Priority and Low Priority Length words indicating the bit lengths of high priority and... ..States Patent No. 5,287,178 discloses a video signal encoding system includes a signal processor for segmenting encoded video data into transport blocks having a header section and a packed data section. The system also includes reset control apparatus for releasing resets of system components, after a global system reset, in a prescribed non-simultaneous phased sequence to enable signal processing to commence in the prescribed sequence. The phased reset release sequence begins when valid data...United States Patent No. 5,142,380 to Sakagami et al. discloses an image compression

apparatus suitable for use with still images such as those formed by electronic still cameras using... ..and Q.

United States Patent No. 5,193,002 to Guichard et al. disclosed an apparatus for coding/decoding image signals in real time in conjunction with the CCITT standard H.261. A digital signal processor carries out direct quantization and reverse quantization.

United States Patent No. 5,241,383 to Chen et al. describes an apparatus with a pseudo-constant bit rate video coding achieved by an adjustable quantization parameter. The... ..relates to an improved pipeline system having an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along the pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the processing stages in the pipeline and interacting with selected stages in the pipeline for control data and/or combined control-data functions among the processing stages, so that the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing. In accordance with the invention, the processing stages may be configurable in response to recognition of at least one token. One of the processing stages may be a Start Code Detector which receives the input and generates and/or converts... ..and resetting the system, and a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards. The present invention... ..data and having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video pictures of varying data size can be controlled. Also in accordance with the invention, a processing stage receives the input data stream, the stage including means for recognizing specified bit

stream patterns, whereby the processing stage facilitates random access and error recovery. The invention may also include a means for... ..invention also includes an inverse modeller stage, an inverse discrete cosine transform stage, and a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for processing data.

In addition, the present invention relates to an improved pipeline system having a Huffman...pipeline stage that incorporates a two-wire transfer control and also shows two consecutive pipeline processing stages with the two-wire transfer control;

Figures. 5a and 5b taken together depict one... ..shown in Figures. 8a and 8b.

Figure 10 is a block diagram of a reconfigurable processing stage;

Figure 11 is a block diagram of a spatial decoder;

Figure 12 is a... ..decoder including the prediction filters;

Figure 18 is a pictorial representation of the prediction filtering process;

Figure 19 shows a generalized representation of the macroblock structure;

Figure 20 shows a generalized... ..buffer;

Figure 25 is a pictorial diagram illustrating prediction data offset from the block being processed;

Figure 26 is a pictorial diagram illustrating prediction data offset by (1,1);

Figure 27...in general terms, the present invention provides an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along a pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the stages in the pipeline and interacting with selected stages in the pipeline for control, data and/or combined control-data functions among the processing stages, whereby the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing.

Each of the processing stages in the pipeline may include both primary and secondary storage, and the stages in... ..The tokens in the pipeline are dynamically adaptive and may be position dependent upon the processing stages for performance of functions or position Independent of the processing stages for performance of functions.

In a pipeline machine, In accordance with the invention, the... ..altered by interfacing with the stages, and the tokens may interact with all of the processing stages in the pipeline or only with some but less than all of said processing stages. The tokens in the pipeline may interact with adjacent processing stages or with non-adjacent processing stages, and the tokens may reconfigure the processing stages. Such tokens may be position dependent for some functions and position independent for other... ..be Huffman coded.

In the improved pipeline machine, the tokens may be generated by a processing stage. Such pipeline tokens may include data for transfer to the processing stages or the tokens may be devoid of data. Some of the tokens may be identified as DATA tokens and provide data to the processing stages in the pipeline, while other tokens are identified as control tokens and only condition the processing stages in the pipeline, such conditioning including reconfiguring of the processing stages. Still other tokens may provide both data and conditioning to the processing stages in the pipeline. Some of said tokens may identify coding standards to the processing stages in the pipeline, whereas other tokens may operate independent of any coding standard among the processing stages. The tokens may be capable of successive alteration by the processing stages in the pipeline.

In accordance with the invention, the interactive flexibility of the tokens in cooperation with the processing stages facilitates greater functional diversity of the processing stages for resident structure in the pipeline, and the flexibility of the tokens facilitates system... ..or alteration. The tokens may be capable of facilitating a plurality of

functions within any processing stage in the pipeline. Such pipeline tokens may be either hardware based or software based... ..system bandwidth in the pipeline. The tokens may provide data and control simultaneously to the processing stages in the pipeline.

The invention may include a pipeline processing machine for handling plurality of separately encoded bit streams arranged as a single serial bit... ..and for passing unrecognized control tokens along the pipeline, and a reconfigurable decode and parser processing means responsive to a recognized control token for reconfiguring a particular stage to handle an... ..be a pipeline system and the Start Code Detector may be positioned as the first processing stage in the pipeline.

The present invention also provides, in a system having a plurality of processing stages, a universal adaptation unit in the form of an interactive interfacing token for control and/or data functions among the processing stages, the token being a PICTURE(underscore)START code token for indicating that the start... ..The token may also be a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards.

The CODING(underscore... ..picture standard as JPEG, and/or any other appropriate picture standard. At least some of the processing stages reconfigure in response to the CODING(underscore)STANDARD token.

One of the processing stages in the system may be a Huffman decoder and parser and, upon receipt of... ..Data stage, and the parser stage may send an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... ..video data, having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video... ..controlled.

The system may include a spatial decoder having a two-wire interface interconnecting processing stages, the interface enabling serial processing for data and parallel processing for control.

As previously indicated, the system may further include a ROM having separate stored... ..each of a plurality of picture standards, the programs being selectable by token to facilitate processing for a plurality of different picture standards.

The spatial decoder system also includes a token... ..decoding stage and a parser stage for sending an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... ..The present invention also provides a pipeline system having an input data stream, and a processing stage for receiving the input data stream, the stage including means for recognizing specified bit... ..whereby said stage facilitates random access and error recovery. In accordance with the invention, the processing stage may be a start code detector and the bit stream patterns may include start... ..token and padding insures uniformity of word size. In accordance with the invention, a reconfigurable processing stage may be provided as a spatial decoder and the padding means adds to picture... ..that if the DATA token has less than the predetermined length, the padder circuit adds units of data to the DATA token until the predetermined length is achieved. A bypass circuit...1 tokens into a buffer, having a second predetermined width.

The invention also provides an apparatus for providing a time delay to a group of compressed pictures, the pictures corresponding to... ..and capable of delaying the words of data, is in communication with a control circuit intermediate the counter circuit and the inverse modeller circuit, the control circuit also communicating with the... ..inverse modeller stage and an inverse discrete cosine transform stage, the improvement characterized by a

processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for

processing data.

In accordance with the invention, the token may be a QUANT(underscore)TABLE token for causing the processing stage to generate a quantization table.

The present invention also provides a Huffman decoder for decoding... ..of bits used to represent an item of data.

DECODER: An embodiment of a decoding process.

DECODING (PROCESS): The process defined in this specification that reads an input coded bitstream and produces decoded pictures or... ..the same order in which they were presented at the input of the encoder.

ENCODING (PROCESS): A process, not specified in this specification, that reads a stream of input pictures or audio samples...to provide an estimate of the pel value or data element currently being decoded.

RECONFIGURABLE PROCESS STAGE (RPS): A stage, which in response to a recognized token, reconfigures itself to perform various operations.

SLICE: A series of macroblocks.

TOKEN: A universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions.

START... ..indicates that the corresponding stage holds valid data, i.e., data that is to be processed in one of the pipeline stages. After processing (which may involve nothing more than a simple transfer without manipulation of the data) valid... ..present invention may be used with any number of pipeline stages. Furthermore, data may be processed in more than one stage and the processing time for different stages can differ.

In addition to clock and data signals (described below... ..other system. For example, the last pipeline stage may pass its data on to subsequent processing circuitry. The ACCEPT signal, which is illustrated as the lower of the two lines connecting... ..the minimum disturbance possible to other pipeline stages. Succeeding pipeline stages are allowed to continue processing and, therefore, this means that gaps open up in the stream of data following the... ..The data in the pipeline is encoded such that many different types of data are processed in the pipeline. This encoding accommodates data packets of variable size and the size of...the other hand, it may generate itself, all or part of the data to be processed in the pipeline. Indeed, as is explained below, a "stage" may contain arbitrary processing circuitry, including none at all (for simple passing of data) or entire systems (for example... ..values zero and 255 may not be used.

If such a picture were to be processed in a pipeline built in the practice of the present invention, then one of these... ..data must not be written over since it is data that must be saved for processing or use in a downstream device e.g., a pipeline stage, a device or a... ..connected to the pipeline upstream contains data D4 that is to be transferred into and processed in the pipeline. Stages B, D and E, in addition to the upstream device, contain...pipeline, in accordance with the preferred embodiments of the present invention, to "fill up" empty processing stages is highly advantageous since the processing stages in the pipeline thereby become decouple from one another. In other words, even though... ..data can be transferred into the pipeline and between stages even when one or more processing stages is blocked.

In the embodiment shown in Fig. 1, it is assumed that the...propagate all the way back to the beginning of the pipeline if there is some intermediate stage that is able to accept new data.

In the embodiment illustrated in Fig. 1...has been mentioned. It is to be further understood that each pipeline stage may also process the data it has received arbitrarily before passing it between its internal storage elements or... ..the portion of the pipeline that contains input and output storage elements and that arbitrarily processes data stored in its storage elements.

Furthermore, the "device" downstream from the pipeline Stage F...valid data, but also when a stage requires more than one clock phase to finish processing its data. This also can occur when it creates valid data in one or both...

...control the passage of data between adjacent storage elements. The VALID signal may also be processed in an analogous manner.

A great advantage of the two-wire interface (one wire for... ..In addition, two extra latches and a small number of gates are preferably added to process the ACCEPT and VALID signals that are associated with the data latches in each half... ..application so requires. The interface in accordance with this embodiment can also be used to process analog signals.

As discussed previously, while other conventional timing arrangements may be used, the interface... ..circuit B1, which may be provided to convert output data from input latch LDIN into intermediate data, which is then later loaded in an output data latch LDOUT, which comprises the... ..is connected either directly as an input to the validation output latch LVOUT, or via intermediate logic devices or circuits that may alter the signal.

Similarly, the output validation signal QVOUT... ..to the input of the validation input latch QVIN of the following stage, or via intermediate devices or logic circuits, which may alter the validation signal. This output QVIN is also...word.

Preferred Data Structure - "tokens"

In the sample application shown in Fig. 4, each stage processes all input data, since there is no control circuitry that excludes any stage from allowing... ..are connected together in a relatively simple configuration. The simplest configuration is a pipeline of processing steps. For example, in the one shown in Fig. 1. The use of tokens, however...flows from left to right in the diagram. Data enters the machine and passes into processing Stage A. This may or may not modify the data and it then passes the... ..advantage of the tokens is their ability to achieve this kind of communication. Since any processing stage that does not recognize a token simply passes it on unaltered to the next... ..is transmitted along with the address and data fields in each token so that a processing stage can pass on a token (which can be of arbitrary length) without having to... ..be the first word of a new token.

Note that although the simple pipeline of processing stages is particularly useful, it will be appreciated that tokens may be applied to more complicated configurations of processing elements. An example of a more complicated processing element is described below.

It is not necessary, in accordance with the present invention, to... ..has extension bits. An example of this is a token that activates a stage that processes video quantization values stored in a quantization table (typically a memory device). For example, a...turn, is of great importance in video data pipeline systems since it ensures that all processing stages can be continuously running at full bandwidth.

In accordance to the present invention, in... ..some other chips in the set This is advantageous both from the perspective of a customer and from that of a chip manufacturer. Even if modifications mean that all chips are...the end of a token (and hence the start of the next token) to be processed correctly (including simple non-manipulative transfer), even if the token is not recognized by the... ..block diagram of a pipeline stage whose function is as follows. If the stage is processing a predetermined token (known in this example as the DATA token), then it will duplicate... ..the address field of the DATA token. If, on the other hand, the stage is processing any other kind of token, it will delete every word. The overall effect is that... ..respective output signals:

In the duplication stage, the output from the data latch LDIN forms intermediate data referred to as MID(underscore)DATA. This intermediate data word is loaded into the data output latch LDOUT only when an intermediate acceptance signal (labeled "MID(underscore)ACCEPT" in Fig. 8a) is set HIGH.

The portion of... ..data. These include a "DATA(underscore)TOKEN" signal that indicates that the circuitry is currently processing a valid DATA Token, and a NOT(underscore)DUPLICATE signal which is used to control duplication of data. When the circuitry is processing a DATA Token, the NOT(underscore)DUPLICATE signal toggles between a HIGH and a LOW... ..the token to be duplicated once (but no more times). When the circuitry is

not processing a valid DATA Token then the NOT(underscore)DUPLICATE signal is held in a HIGH state. Accordingly, this means that the token words that are being processed are not duplicated.

As Fig. 8a illustrates, the upper six bits of 8-bit intermediate data word and the output signal all from the latch LI1 form inputs to a... ..explained further below.

Latch LO1 performs the function of latching the last value of the intermediate extension bit (labeled "MID(underscore)EXTN" and as signal S4), and it loads this value... ..and the DATA(underscore)TOKEN signal will become "0", indicating that the circuitry is not processing a DATA token.

If QI1 is "0" and SO is "0", thereby indicating a DATA... ..phase and the DATA(underscore)TOKEN signal will become "1", indicating that the circuitry is processing a DATA token.

The NOT(underscore)DUPLICATE signal (the output signal QO3) is similarly loaded...LVOUT at the same time that MID(underscore)DATA is loaded into LDOUT and the intermediate extension bit (signal S4) is loaded into LEOUT. Signal S5 is also combined with the...above. This has the effect that all tokens except the one that causes the duplication process will be deleted from the token stream, since a device connected to the output terminals (OUTDATA, OUTEXTN and OUTVALID) will not recognize these token words as valid data.

As before... ..and is duplicated.

Referring now more particularly to Figure 10, there is shown a reconfigurable

process stage in accordance with one aspect of the present invention.

Input latches 34 receive an... ..the input latches 34 is passed as a first input over line 35 to a

processing unit 36. A first output from the token decode subsystem 33 is passed over line 37 as a second input to the processing unit 36. A second output from the token decode 33 is passed over line 40 to an action identification unit 39. The action identification unit 39 also receives input from registers 43 and 44 over line 46. The registers 43... ..is determined by the history of tokens previously received. The output from the action identification unit 39 is passed over line 38 as a third input to the processing unit 36. The output from the processing unit 36 is passed to output latches 41. The output from the output latches 41 is... ..decoder 56 is passed over line 63 as an input to an Index to Data Unit (ITOD) 64. The Huffman decoder 56 and the ITOD 64 work together as a single logical unit. The output from the ITOD 64 is passed over line 65 to an arithmetic logic unit (ALU) 66. A first output from the ALU 66 is passed over line 67 to...blocks 133.

Referring to Figure 14b, in the JPEG and H.261 standards, the Common Intermediate Format(CIF) is used, wherein a picture 141 is encoded as 6 rows each containing... ..in a zigzag direction indicated by the arrow 144. The GOBs 142 are, in turn, processed row-by-row, left-to-right in each row.

Referring now to Figure 14c, it... ..in accordance with the practice of the present invention. A first picture 161 to be processed contains a first PICTURE(underscore)START token 162, first-picture information of indeterminate length 163, and a first PICTURE(underscore)END token 164. A second picture 165 to be processed contains a second PICTURE(underscore)START token 166, second picture information of indeterminate length 167... ..tokens 162 and 166 indicate the start of the pictures 161 and 165 to the processor. Likewise, the PICTURE(underscore)END tokens 164 and 168 signify the end of the pictures 161 and 165 to the processor. This allows the processor to process picture information 163 and 167 of variable lengths.

Referring to Figure 17, a split 171...Video Formatter (not shown in Figure 17).

Referring now to Figure 18, the prediction filtering process is illustrated. A forward picture 201 is passed over line 202 as a first input... ..the right of the value decode shift register 230, as indicated by area 231. This process

eliminates overlapping start code images, as discussed below. A first output from the value decode... ..Code Detector. The Start Code Detector then receives a first data value image 244. Before processing the first data value image 244, the Start Code Detector may detect a second start... ..image 244 at a length 246. If this occurs, the Start Code Detector does not process the first data value image 244, and instead receives and processes a second data value image 247.

Referring now to Figure 22, a flag generator 251...line 1 of Table 600, whenever a "sequence start" image is received during H.261 processing or a "picture start" image is received during MPEG processing, the entire group of four control tokens is generated, each followed by its corresponding data...Picture Decoding

3. Motion Picture Decompression

4. RAM Memory Map

5. Bitstream Characteristics

6. Reconfigurable Processing Stage

7. Multi-Standard Coding

8. Multi-Standard Processing Circuit-2nd Mode of Operation

9. Start Code Detector

10. Tokens

11. DRAM Interface

12... ..described herein in greater detail) and reformatting this output for use, including display in a computer or other display systems, including a video display system. Implementation of this formatting varies significantly... ..the Spatial Decoder circuits.

The Spatial Decoder of the present invention performs all the required processing within a single picture. This reduces the redundancy within one picture.

The Temporal Decoder reduces... ..modeller 75, the inverse zig-zag 81 and the inverse DCT 83. The standard independent units within the Huffman decoder and parser include the ALU 66 and the token formatter 71.

Referring now to Figure 12, the standard-independent units include the DRAM interface 100, the fork 91, the FIFO register 96, the summer 98 and the output selector 106. The standard dependent units are the address generator 94, which is different in H.261 and in MPEG, and...much of the operation is very similar between the three different compression standards.

The next unit is the state machine 68 (Figure 11) located within the Huffman decoder and parser. Here... ..The same holds true for JPEG, which is a third, completely independent program.

The next unit is the Huffman decoder 56 which functions with the index to data unit 64. Those two units cooperate together to perform the Huffman decoding. Here, the algorithm that is used for Huffman... ..to the Huffman decoder at different times consistent with the standard in operation.

The last unit on the chip that is dependent on the compression standard is the inverse quantizer 79...an H.261 group of blocks and an MPEG slice. When H.261 data is processed after the Start Code Detector, each group of blocks is preceded by a slice(underscore... ..these standards have totally different sets of tables.

As previously indicated, most of the system units are compression standard independent. If a unit is standard independent, and such units need not remember what CODING(underscore)STANDARD is being processed. All of the units that are standard dependent remember the compression standard as the CODING(underscore)STANDARD

token flows... ..CODING(underscore)STANDARD tokens at the Start Code Detector that is positioned as the first unit in the pipeline, this change of compression standard is readily handled. The token says a... ..found in the standard, i.e. from the bitstream into a prediction mode token. This processing is performed by the Huffman decoder and parser state machine, where it is easy to... ..to that token. By having these tokens and using them appropriately, the design of other units in the machine is simplified. Although there may be some complications in the program, benefits... ..a first encoded signal (the MPEG or H.261 encoded video signal) in a pipeline processing system. The Temporal Decoder is not needed for JPEG decoding.

In this regard, the invention... ..the use of a single pipeline decoder and decompression system. The decoding and decompression pipeline processor is organized on a unique and special configuration which allows the handling of the multi... ..video signals through the use of techniques all compatible with the single pipeline decoder and processing system. The Spatial Decoder is combined with the Temporal Decoder, and the Video Formatter is used... ..with only still pictures. The compression standard independent Spatial Decoder performs all of the data processing within the boundaries of a single picture. Such a ...to the multi-standard, configurable Video Formatter, which then provides an output to the display terminal. In a first sequence of similar pictures, each decompressed picture at the output of the... ..of control tokens and DATA tokens, in combination with a plurality of sequentially-positioned reconfigurable processing stages selected and organized to act as a standard independent, reconfigurable-pipeline-processor.

With regard to JPEG decoding, a single Spatial Decoder with no off chip DRAM can... ..video. Accordingly, signals carried by DATA tokens pass directly through the Temporal Decoder without further processing when the Temporal Decoder is configured for a JPEG operation.

Another aspect of the present... ..for subsequent use in temporal decoding of subsequent pictures.

Generally, the Temporal Decoder performs the processing between pictures either earlier and/or later in time with reference to the picture currently being... ..is distributed among several areas of DRAM in the sense that the decompressed output information, processed by the Spatial Decoder, is stored in other DRAM registers by other random access memories...first decoder circuit (the Spatial Decoder) directly to the Video Formatter for handling without signal processing delay.

The Temporal Decoder also reorders the blocks of picture data for display by a... ..from a selection of pictures which have arrived earlier or later than the picture under processing. When a picture is described in this context, it may mean any one of the... ..2. The result, i.e., the final decoded picture resulting from the addition of a process step performed by the decoder;

3. Previously decoded pictures read from the DRAM; and

4... ..START token and a subsequent PICTURE(underscore)END token.

After the picture data information is processed by the Temporal Decoder, it is either displayed or written back into a picture memory location. This information is then kept for further reference to be used in processing another different coded data picture.

Re-ordering of the MPEG encoded pictures for visual display...used to encode a referenced picture of a picture might be identified as being one unit long, another picture might be a number of units long, while still a third picture could be a fraction of that unit

None of the existing standards (MPEG 1.2, JPEG, H.261) define a way of... ..picture rate, whereas the Video Formatter can handle a variable input picture rate.

6. RECONFIGURABLE PROCESSING STAGE

Referring again to Figure 10, the reconfigurable processing stage (RPS) comprises a token decode circuit 33 which is employed to receive the tokens... input latches 34. The output of the token decode circuit 33 is applied to a processing unit 36 over the two-wire interface 37 and an action identification circuit 39. The processing unit 36 is suitable for processing data under the control of the action identification circuit 39. After the processing is completed, the

processing unit 36 connects such completed signals to the output, two-wire interface bus 40 through output... token decode circuit 33 are applied simultaneously to the action identification circuit 39 and the

processing unit 36. The action identification function as well as the RPS is described in further detail... not standard independent circuits. The data flows through the token decode circuit 33, through the processing unit 36 and onto the two-wire interface circuit 42 through the output latches 41. If... wire interface 42 through the output circuit 41. The present invention operates as a pipeline processor having a two-wire interface for controlling the movement of control tokens through the pipeline... time, the token decode circuit 33 provides a proper flag or index signal to the processing unit 36 to alert it to the presence of the token being handled by the action identification circuit 39.

Control tokens may also be processed.

A more detailed description of the various types of tokens usable in the present invention passing through the state machine shown with reference to Figure 10.

Similarly, the processing unit 36 which is under the control of the action identification circuit 39 is now ready to process the information contained in the data fields of the DATA token when it is appropriate... action identification circuit 39 and is immediately followed by a DATA token which is then processed by the processing unit 36. The control token exits the output latches circuit 41 over the output two-wire Interface 42 immediately preceding the DATA token which has been processed within the processing unit 36.

In the present invention, the action identification circuit, 39, is a state machine holding... show that the action can also be affected by the token that is currently being processed by the token decode circuit 33.

In general, there is shown token decoding and data processing in accordance with the present invention. The data processing is performed as configured by the action identification circuit 39. The action is affected by... information stored from previously decoded tokens in registers 43 and 44, the current token under processing, and the state and history information that the action identification unit 39 has itself acquired. A distinction is thereby shown between Control tokens and DATA tokens.

In any RPS, some tokens are viewed by that RPS unit as being Control tokens in that they affect the operation of the RPS presumably at... are viewed by the RPS as DATA tokens. Such DATA tokens contain information which is processed by the RPS in a way that is determined by the design of the particular... view of the same token. Some of the tokens might be viewed by one RPS unit as DATA Tokens while another RPS unit might decide that it is actually a Control Token. For example, the quantization table information... into a token-called a quantization table token (QUANT(TABLE)) which goes down the processing pipeline. As far as that machine is concerned, all of that was data; it was... sort of data into another sort of data, which is clearly a function of the processing performed by that portion of the machine. However, when that information gets to the inverse... present. This information is viewed as control information, and then that control information affects the processing that is done on subsequent DATA tokens because it affects the number that you multiply... important feature of the invention is that each of the stages of circuitry has the processing capability within it to be able to perform the necessary operations for each of the... operations are to be performed at a given time; come as tokens. There is one processing element that differs between the different stages to provide this capability. In the state machine... standard is and it looks up the parameters that it needs to apply to the processing elements in order to perform a proper

operation. For example, the inverse quantizer will look... ..is set to 1 for a particular compression standard, and will apply that to its processing circuitry.

In a similar sense the Huffman decoder 56 has a number of tables within... ..MPEG video standard or the JPEG video standard. These three compression coding standards specify similar processes to be done on the arriving data, but the structure of the datastreams is different... ..token stream embodying the current coding standard. The control tokens are passed through the pipeline processor, and are used, i.e., decoded, in the state machines to which they are relevant... ..this regard, the DATA Tokens are treated in the same fashion, insofar as they are processed only in the state machines that are configurable by the control tokens into processing such DATA Tokens. In the remaining state machines, they pass through unchanged.

More specifically, a... ..signals. The remaining portions of the token are used to indicate and identify the internal processing control function which is standard for all of the datastreams passing through the pipeline processor. In one form of the invention, the token extension is used to carry the current... ..accompanying data. As previously discussed, this information is utilized in the system to reconfigure the processing stage used to perform the function required by the various standards created for that purpose... ..picture number as indicated by the value.

The system also includes a multi-stage parallel processing pipeline operating under the principles of the two-wire interface previously described. Each of the... ..the token presently entering the state machine into the action identification circuit 39 or the processing unit 36, as appropriate. The processing unit has been previously reconfigured by the next previous control token into the form needed for handling the current coding standard, which is now entering the processing stage and carried by the next DATA token. Further, in accordance with this aspect of the invention, the succeeding state machines in the processing pipeline can be functioning under one coding standard, i.e., H.261, while a previous... ..tokens required to decode a number of coding standards with a fixed number of reconfigurable processing stages. More specifically, the PICTURE(underscore)END control token is employed because it is important... ..standard machine, it is necessary to create additional control tokens within the multi-standard pipeline processing machine which will then indicate which one of the standard decoding techniques to use. Such... ..and to push the current picture through the decoder to the display.

8. MULTI-STANDARD PROCESSING CIRCUIT - SECOND MODE OF OPERATION

A compression standard-dependent circuit, in the form of the... ..of the Start Code Detector will subsequently be discussed in further detail, as will the process of starting up of the decoder.

The aforementioned description has been concerned primarily with the... ..the data which immediately follows according to the standard. However, in the multi-standard pipeline processing system of the present invention, where compatibility is required for multiple standards, the system has... ..signals, including flag signals, are generated by each state machine to handle some of the processing within that state machine. Values carried in the standards can be used to access machine... ..its contents must be removed from the two wire interface to ensure that no further processing takes place using these 3 bytes. The decode register is emptied, and the value decode... ..10. TOKENS

In the practice of the present invention, a token is a universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions and is adapted for use with a reconfigurable processing stage (RPS) which is a stage, which in response to a recognized token, reconfigures itself to perform various operations.

Tokens may be either position dependent or position independent upon the processing stages for performance of various functions. Tokens may also be metamorphic in that they can be altered by a processing stage and then passed down the pipeline for performance of further functions. Tokens may interact... ..other functions, and the specific interaction with a stage may be conditioned by the previous processing history of a stage.

A PICTURE(underscore)END token is a way of signalling the... ..through a fixed size, fixed width buffer.

The present invention is directed to a pipeline processing system which has a variable configuration which uses tokens and a two-wire system. The... ..do not use control tokens.

The control tokens are generated by circuitry within the decoder processor and emulate the operation of a number of different type standard-dependent signals passing into the serial pipeline processor for handling. The technique used is to study all the parameters of the multi-standards that are selected for processing by - the serial processor and noting 1) their similarities, 2) their dissimilarities, 3) their needs and requirements and 4) selecting the correct token function to effectively process all of the standard signals sent into the serial processor. The functions of the tokens are to emulate the standards. A control token function is... ..the standard dependent signals and as an element to transmit control information through the pipeline processor.

In prior art system, a dedicated machine is designed according to well-known techniques to... ..tokens provide and make a sensible format for communicating information through the decompression circuit pipeline processor. In the design selected hereinafter and used in the preferred embodiment, each word of a... ..However, this is not a limitation on the invention, but on the magnitude of the processing steps elected to be accomplished by use of these tokens. It is to be noted... ..bit address for use in accessing the random access memories used throughout this serial decompression

processor. This provides an additional degree of variability that facilitates a broad range of versatility.

As previously described, the DATAtoken carries data from one

processing stage to the next. Consequently, the characteristics of this token change as it passes through...longest number of data bits because it needs to provide the most information to the processing unit so that it can start the decompression with as much information as possible. Words which...to receive an address, it waits for the address generator to supply a valid address, processes that address and then sets the accept line high for one clock period. Thus, it... ..be read. This signal passes between two asynchronous clock regimes and, therefore, passes through three synchronizing flip flops.

Provided RAM2 312 is empty, the next item of data to arrive on...interesting.

In general, prediction data will be offset from the position of the block being processed as specified in the motion vectors in x and y. Thus, the Mock of data... ..address, 9. Data is read from this address and the x value is incremented. The process is repeated until the x value reaches its stop value, at which point, the y... ..is read, the x value is again incremented until it reaches its stop value. The process is repeated until both x and y values have reached their stop values. Thus, the...invention, is that additional information must be provided to the prediction filters to indicate what processing is required on the data. This consists of the following:

a last byte" signal indicating... ..bit 0) is incremented and the x address (3 LSBS) is reset to zero. This process is repeated until 64 bytes have been read. With a 16 or 32 bit wide...register while its access register is set to zero, the results are undefined.

14. MICRO-PROCESSOR INTERFACE

A standard byte wide micro-processor interface (MPI) is used on all circuits with in the Spatial Decoder and Temporal Decoder... ..the parameter column. The actual specifications are shown in the respective columns min, max and units.

The DC operating conditions can be seen with reference to Table A.6.3. Here... ..the signal is present the maximum amount of time that this signal is available. The Units column gives the units of measurement used to describe the signals.

16. MPI WRITE TIMING

The general description of...Consequently, the machine will not go into error recovery mode and will successfully continue to process the coded data.

A still further advantage of the use of a PICTURE(underscore)END token is that the serial pipeline processor will continue the processing of uninterrupted data. Through the use of a PICTURE(underscore)END token, the serial pipeline processor is configured to handle less than the expected amount of data and, therefore, continues processing. Typically, a prior art machine would stop itself because of an error condition. As previously... ..of the Huffman decode and Video Demultiplexor know the number of blocks that it will process during each picture recovery cycle. When the correct number of blocks do not arrive from the... ..Each of the state machines recognizes a FLUSH control token as information not to be processed. Accordingly, the FLUSH token is used to fill up all of the remaining empty parts...less information than normally expected to decode the last picture. The Huffman decode circuit finishes processing the information contained in the last picture, and outputs this information through the DRAM interface... ..token, in accordance with the present invention, is used to pass through the entire pipeline processor and to ensure that the buffers are emptied and that other circuits are reconfigured to... ..underscore)END token, a padding word and a FLUSH token indicating to the serial pipeline processor that the picture processing for the current picture form is completed. Thereafter, the various state machines need reconfiguring to... ..FLUSH token resets each stage as it passes through, but-allows subsequent stages to continue processing. This prevents a loss of data. In other words, the FLUSH token is a variable... ..AFTER PICTURE

The STOP(underscore)AFTER(underscore)PICTURE function is employed to shut down the processing of the serial pipeline decompressing circuit at a logical point in its operation. At this... ..a picture, the STOP(underscore)AFTER(underscore)PICTURE operation signals the end of all current processing.

22. MULTI(underscore)STANDARD - SEARCH MODE

Another feature of the present invention is the use... ..underscore)MODE control token which is used to reconfigure the input to the serial pipeline processor to look at the incoming bit stream. When the search mode is set, the Start... ..combination of control tokens, and DATA tokens along with the reconfiguration circuits, to provide similar processing.

The use of search mode in the present invention is convenient in many situations including... ..video disc. In general, a search mode is convenient when the user interrupts the normal processing of the serial pipeline at a point where the machine does not expect such an...be the case.

In brief, the Huffman Decoder 321 works in conjunction with the other units shown in Figure 27. These other units are the Parser State Machine 322, the inshifter 323, the Index to Data unit 324, the ALU 325, and the Token Formatter 326. As described previously, connection between these blocks is governed by a two wire interface. A more detailed description of how these units function is subsequently described herein in greater detail, the focus here is on particular aspects... ..control certain functions of the Index to Data 324 and ALU 325. Control of these units by the Huffman Decoder is necessary for proper decoding of block-level information. Having the... ..further detail in the "More Detailed Description of the Invention" section.

The Index to Data unit 324 performs the second part of the multi-part algorithm. This unit contains a look up table that provides the actual Huffman decoded data. Entries in the...by detecting these in the Huffman Decoder 321, rather than in the Index to Data unit 324.

This index number is then passed to the Index to Data unit 324. In essence, the Index to Data unit is a look-up table. In accordance with one aspect of the algorithm, the look... ..that J PEG specifies for transferring an alternate JPEG table.

From the Index to Data unit 324, the decoded index number or other data is passed, together with the accompanying control... ..the entering data to ensure that the DATA tokens are of the correct size for processing. In fact, the token stream can be corrected in some situations if the error is... ..an order that is useful for the decompression circuits, but

not for the particular display unit being used. When a block of data enters the Buffer Manager, the Buffer Manager supplies...the output of the Spatial Decoder or Temporal Decoder and re-format it for a computer or display system. The details of this formatting will vary between applications. In a simple...Token. The DATA Token can have as many bits as are necessary for carrying out processing at a particular place in the system. All other Tokens ignore the extra bits.

A.3.2 The DATA Token

The DATA Token carries data from one processing stage to the next. Consequently, the characteristics of this Token change as it passes through... will be sufficient to collect DATA Tokens and to detect a few Tokens that provide synchronization information (such as PICTURE(underscore)START). In this regard, see subsequent sections A.16, "Connecting... from the data stream. This provides an alternative to doing the configuration via the micro processor interface.

A.3.4 Description of Tokens

This section documents the Tokens which are implemented... 3.5.1. Note: JPEG requires a 2:1:1 structure for its macroblocks when processing 4:2:2 data. See Table A.3.5.

A.3.6 Special Token formats...either is low then the interface is taken to high impedance.

Note: on-chip data processing is not terminated when the DRAM interface is at high impedance. Therefore, errors will occur...decoded video's picture rate. Accordingly, this clock can be used to provide audio/video synchronization.

A.7.1 Spatial Decoder clock signals

The Spatial Decoder has two different (and potentially...in accordance with the present invention, must know what video standard is being input for processing. Thereafter, the system can accept either pre-existing Tokens or raw byte data which is... time a value is written into coded(underscore)data (7:0). Software is responsible for settling coded(underscore)extn to 0 before the last word of any Token is written to... 0). The start of this new DATA Token then passes into the Spatial Decoder for processing.

Each time a new 8 bit value is written to coded(underscore)data (7:0... Detector analyses data in the DATA Tokens bit serially. The Detector's normal rate of processing is one bit per clock cycle (of coded(underscore)clock). Accordingly, it will typically decode a byte of coded data every 8 cycles of coded(underscore)clock. However, extra processing cycles are occasionally required, e.g., when a non-DATA Token is supplied or when... Furthermore, this clock can be asynchronous to the main decoder(underscore)clock Data transfer is synchronized to decoder(underscore)clock on-chip.

SECTION A.11 Start code detector

A.11.1... Code Detector. So, accessing these registers will be unreliable if the Start Code Detector is processing data. The user is responsible for ensuring that the Start Code Detector is halted before... Detector. In this case, the Tokens are passed through the Start Code Detector with no processing to other stages of the Spatial Decoder. These Tokens can only be inserted just before...result will be unpredictable if this is done when the Start Code Detector is actively processing data. Discard all mode can be safely initiated after any of the Start Code Detector... start code non-alignment interrupt is suppressed.

In contrast, however, JPEG was designed for a computer environment where byte alignment is guaranteed. Therefore, marker codes should only be detected when byte... the other hand, was designed to meet the needs of both communications (bit serial) and

computer (byte oriented) systems. Start codes in MPEG data should normally be byte aligned. However, the...result will be unpredictable if this is done when the Start Code Detector is actively processing data. So, before initiating a start code search, the Start Code Detector should be stopped so no data is being processed. The Start Code Detector is always in this condition if any of the Start Code...

Claims:

1. A pipeline processing machine having a plurality of reconfigurable processing stages interconnected by a two-wire interface bus, one of said processing stages being a spatial decoder; a second of said stages being a token generator for...

6/K/13 (Item 13 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Country	Number	Kind	Date		
Type		Pub. Date		Kind	Text
Available Text		Language		Update	Word Count
Total Word Count (Document A)					
Total Word Count (Document B)					
Total Word Count (All Documents)					

Specification:

The present invention is directed to improvements in methods and apparatus for decompression which operates to decompress and/or decode a plurality of differently encoded input... ..of the well known standards known as JPEG, MPEG and H.261.

A serial pipeline processing system of the present invention comprises a single two-wire bus used for carrying unique... ..to a plurality of adaptive decompression circuits and the like positioned as a reconfigurable pipeline processor,

PRIOR ART

One prior art system is described in United States Patent No. 5,216,724. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and a host processor. The device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus.

United States Patent No. 4,785... ..a known quad tree data structure.

United States Patent No. 5,122,875 discloses an apparatus for encoding/decoding an HDTV signal. The apparatus includes a compression circuit responsive to high definition video source signals for providing hierarchically layered... ..compressed video data of relatively greater and lesser importance to image reproduction respectively. A transport processor, responsive to the high and low priority codeword sequences, forms high and low priority transport... ..United States Patent No. 5,168,356 discloses a video signal encoding system that includes apparatus for segmenting encoded video data into transport blocks for signal transmission. The ...in respective transport blocks.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening. This is accomplished by an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ..frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ..oscillators and the receiver can continuously receive each channel, then the receiver need not be synchronized with the transmitter. An FFT algorithm implements a fast discrete approximation to the continuous case in which the receiver synchronizes to the first frame and then acquires subsequent frames every frame period. The frame period... ..increasing the amount of data transmitted.

United States Patent No. 5,212,742 discloses an apparatus and method for processing video data for com(underscore) pression/decompression in real-time. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and host processor. Lastly, the device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus. The method handles assigning portions of the image for each of the processors to operate upon.

United States Patent No. 5,231,484 discloses a system and method... ..MPEG standards. Included are three cooperating components or subsystems that operate to variously adaptively pre-process the incoming digital motion video sequences, allocate bits to the pictures in a sequence, and... ..States Patent No. 5,267,334 discloses a method of removing frame redundancy in a computer system for a sequence of moving images. The method comprises detecting a first scene change... ..facing" keyframe or intraframe, and it is normally present in CCITT compressed video data. The process then comprises generating at least one intermediate compressed frame, the at least one intermediate compressed frame containing difference information from the first image for at least one image following...change, known as a "backward-facing" keyframe. The first keyframe and the at least one intermediate compressed frame are linked for forward play, and the second keyframe and the intermediate compressed frames are linked in reverse for reverse play. The intraframe may also be used... ..of complete scene information.

United States Patent No. 5,276,513 discloses a first circuit apparatus, comprising a given number of prior-art image-pyramid stages, together with a second circuit apparatus, comprising the same given number of novel motion-vector stages, perform cost-effective hierarchical motion analysis (HMA) in real-time, with minimum system processing delay and/or employing minimum system processing delay and/or employing minimum hardware structure. Specifically, the first and second circuit apparatus, in response to relatively high-resolution image data from an ongoing input series of successive... ..a relatively high frame rate (e.g., 30 frames per second), derives, after a certain processing-system delay, an ongoing output series of successive given pixel-density vector-data frames that... ..of successive image frames.

United States Patent No. 5,283,646 discloses a method and apparatus for enabling a real-time video encoding system to accurately deliver the desired number of... ..desired bit allocations.

The article, Chong, Yong M., A Data-Flow Architecture for Digital Image Processing, Wescon Technical Papers: No. 2 Oct./Nov. 1984, discloses a real-time signal processing system specifically designed for image processing. More particularly, a token based data-flow architecture is disclosed wherein the tokens are of... ..width having a fixed width address field. The system contains a plurality of identical flow processors connected in a ring fashion. The tokens contain a data field, a control field and a tag. The tag field of the token is further broken down into a processor address field and an identifier field. The processor address field is used to direct the tokens to the correct data-flow processor, and the identifier field is used to label the data such that the data-flow processor knows what to do with the data. In this way, the identifier field acts as an instruction for the data-flow processor. The system directs each token to a specific data-flow processor using a module number (MN). If the MN matches the MN of the

particular stage... ..to locate the decoder in the preceding stage in order to pre-decode complex decoding processing and to alleviate critical path problems in the logic circuit. The elastic nature of the...of block signal in most cases.

United States Patent No. 4,903,018 discloses a process and data processing system for compressing and expanding structurally associated multiple data sequences. The process is particular to data sets in which an analysis is made of the structure in... ..data series on the basis of the order number of these data elements. The data processing system for performing the processes includes a storage matrix (26) and an index storage (28) having line addresses of the... ..the final actual video.

United States Patent No. 5,080,242 discloses an image signal processing system DPCM encodes the signal, then Huffman and run length encodes the signal to produce... ..tightly packed without gaps for efficient transmission without loss of any data. The tightly packed apparatus has a barrel shifter with its shift modulus controlled by an accumulator receiving code word... ..OR gate is connected to the shifter, while a register is connected to the gate. Apparatus for processing a tightly packed and decorrelated digital signal has a barrel shifter and accumulator for unpacking... ..an inverse DCPM decoder.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening is accomplished by use of an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ..frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ..kernel matrix.

United States Patent No. 5,231,486 discloses a high definition video system processes a bitstream including high and low priority variable length coded Data words. The coded Data... ..packed High Priority Data and packed Low Priority Data by means of respective data packing units. The coded Data is continuously applied to both packing units. High Priority and Low Priority Length words indicating the bit lengths of high priority and... ..States Patent No. 5,287,178 discloses a video signal encoding system includes a signal processor for segmenting encoded video data into transport blocks having a header section and a packed data section. The system also includes reset control apparatus for releasing resets of system components, after a global system reset, in a prescribed non-simultaneous phased sequence to enable signal processing to commence in the prescribed sequence. The phased reset release sequence begins when valid data...United States Patent No. 5,142,380 to Sakagami et al. discloses an image compression

apparatus suitable for use with still images such as those formed by electronic still cameras using... ..and Q.

United States Patent No. 5,193,002 to Guichard et al. disclosed an apparatus for coding/decoding image signals in real time in conjunction with the CCITT standard H.261. A digital signal processor carries out direct quantization and reverse quantization.

United States Patent No. 5,241,383 to Chen et al. describes an apparatus with a pseudo-constant bit rate video coding achieved by an adjustable quantization parameter. The... ..relates to an improved pipeline system having an input, an output and a plurality of processing stages between the input and the output the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along the pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the processing stages in the pipeline and interacting with selected stages in the pipeline for control data and/or combined control-data functions among the processing stages, so that the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing. In accordance with the invention, the processing stages may be configurable in response to recognition of at least one token. One of the processing stages may be a Start Code Detector which

receives the input and generates and/or... ..and resetting the system, and a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards. The present invention... ..data and having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video pictures of varying data size can be controlled. Also in accordance with the invention, a processing stage receives the input data stream, the stage including means for recognizing specified bit stream patterns, whereby the processing stage facilitates random access and error recovery. The invention may also include a means for... ..invention also includes an inverse modeller stage, an inverse discrete cosine transform stage, and a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for processing data.

In addition, the present invention relates to an improved pipeline system having a Huffman...pipeline stage that incorporates a two-wire transfer control and also shows two consecutive pipeline processing stages with the two-wire transfer control;

Figures. 5a and 5b taken together depict one... ..Shown in Figures. 8a and 8b.

Figure 10 is a block diagram of a reconfigurable processing stage;

Figure 11 is a block diagram of a spatial decoder;

Figure 12 is a... ..decoder including the prediction filters;

Figure 18 is a pictorial representation of the prediction filtering process;

Figure 19 shows a generalized representation of the macroblock structure;

Figure 20 shows a generalized... ..buffer;

Figure 25 is a pictorial diagram illustrating prediction data offset from the block being processed;

Figure 28 is a pictorial diagram illustrating prediction data offset by (1,1);

Figure 27...in general terms, the present invention provides an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along a pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the stages in the pipeline and interacting with selected stages in the pipeline for control, data and/or combined control-data functions among the processing stages, whereby the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing.

Each of the processing stages in the pipeline may include both primary and secondary storage, and the stages in the processing stages for performance of functions or position independent of the processing stages for performance of functions.

In a pipeline machine, in accordance with the invention, the... ..altered by interfacing with the stages, and the tokens may interact with all of the processing stages in the pipeline or only with some but less than all of said processing stages. The tokens in the pipeline may interact with adjacent processing stages or with non-adjacent processing stages, and the tokens may reconfigure the processing stages. Such tokens may be position dependent for some functions and position independent for other... ..be Huffman coded.

In the improved pipeline machine, the tokens may be generated by a processing stage. Such pipeline tokens may include data for transfer to the processing stages or the tokens may be devoid of data. Some of the tokens may be identified as DATA tokens and provide data to the processing stages in the pipeline, while other tokens are identified as control tokens and only condition the processing stages in the pipeline, such conditioning including reconfiguring of the processing stages. Still other tokens may provide both data and conditioning to the processing

stages in the pipeline. Some of said tokens may identify coding standards to the processing stages in the pipeline, whereas other tokens may operate independent of any coding standard among the processing stages. The tokens may be capable of successive alteration by the processing stages in the pipeline.

In accordance with the invention, the interactive flexibility of the tokens in cooperation with the processing stages facilitates greater functional diversity of the processing stages for resident structure in the pipeline, and the flexibility of the tokens facilitates system... or alteration. The tokens may be capable of facilitating a plurality of functions within any processing stage in the pipeline. Such pipeline tokens may be either hardware based or software based... system bandwidth in the pipeline. The tokens may provide data and control simultaneously to the processing stages in the pipeline.

The invention may include a pipeline processing machine for handling plurality of separately encoded bit streams arranged as a single serial bit... and for passing unrecognized control tokens along the pipeline, and a reconfigurable decode and parser processing means responsive to a recognized control token for reconfiguring a particular stage to handle an... be a pipeline system and the Start Code Detector may be positioned as the first processing stage in the pipeline.

The present invention also provides, in a system having a plurality of processing stages, a universal adaptation unit in the form of an interactive interfacing token for control and/or data functions among the processing stages, the token being a PICTURE(underscore)START code token for indicating that the start... The token may also be a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards.

The CODING(underscore... standard as JPEG, and/or any other appropriate picture standard. At least some of the processing stages reconfigure in response to the CODING(underscore)STANDARD token.

One of the processing stages in the system may be a Huffman decoder and parser and, upon receipt of... Data stage, and the parser stage may send an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... video data, having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video... controlled.

The system may include a spatial decoder having a two-wire interface interconnecting processing stages, the interface enabling serial processing for data and parallel processing for control.

As previously indicated, the system may further include a ROM having separate stored... of a plurality of picture standards, the programs being selectable by a token to facilitate processing for a plurality of different picture standards.

The spatial decoder system also includes a token... decoding stage and a parser stage for sending an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... The present invention also provides a pipeline system having an input data stream, and a processing stage for receiving the input data stream, the stage including means for recognizing specified bit... whereby said stage facilitates random access and error recovery. In accordance with the invention, the processing stage may be a start code detector and the bit stream patterns may include start... token and padding insures uniformity of word size. In accordance with the invention, a reconfigurable processing stage may be provided as a spatial decoder and the padding means adds to picture... that if the DATA token has less than the predetermined length, the padder circuit adds units of data to the DATA token until the predetermined length is achieved. A bypass circuit... 1 tokens into a buffer, having a second predetermined width.

The invention also provides an apparatus for providing a time delay to a group of compressed pictures, the pictures corresponding to... ..and capable of delaying the words of data, is in communication with a control circuit intermediate the counter circuit and the inverse modeller circuit, the control circuit also communicating with the... ..inverse modeller stage and an inverse discrete cosine transform stage, the improvement characterized by a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for

processing data.

In accordance with the invention, the token may be a QUANT(underscore)TABLE token for causing the processing stage to generate a quantization table.

The present invention also provides a Huffman decoder for... ..of bits used to represent an item of data.

DECODER: An embodiment of a decoding process.

DECODING (PROCESS): The process defined in this specification that reads an input coded bitstream and produces decoded pictures or... ..the same order in which they were presented at the input of the encoder.

ENCODING (PROCESS): A process, not specified in this specification, that reads a stream of input pictures or audio samples...to provide an estimate of the pel value or data element currently being decoded.

RECONFIGURABLE PROCESS STAGE (RPS): A stage, which in response to a recognized token, reconfigures itself to perform various operations.

SLICE: A series of macroblocks.

TOKEN: A universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions... ..indicates that the corresponding stage holds valid data. i.e., data that is to be processed in one of the pipeline stages. After processing (which may involve nothing more than a simple transfer without manipulation of the data) valid... ..present invention may be used with any number of pipeline stages. Furthermore, data may be processed in more than one stage and the processing time for different stages can differ.

In addition to clock and data signals (described below... ..other system. For example, the last pipeline stage may pass its data on to subsequent processing circuitry. The ACCEPT signal, which is illustrated as the lower of the two lines connecting... ..the minimum disturbance possible to other pipeline stages. Succeeding pipeline stages are allowed to continue processing and, therefore, this means that gaps open up in the stream of data following the... ..The data in the pipeline is encoded such that many different types of data are processed in the pipeline. This encoding accommodates data packets of variable size and the size of...the other hand, it may generate itself, all or part of the data to be processed in the pipeline. Indeed, as is explained below, a "stage" may contain arbitrary processing circuitry, including none at all (for simple passing of data) or entire systems (for example... ..values zero and 255 may not be used.

If such a picture were to be processed in a pipeline built in the practice of the present invention, then one of these... ..data must not be written over since it is data that must be saved for processing or use in a downstream device e.g., a pipeline stage, a device or a... ..connected to the pipeline upstream contains data D4 that is to be transferred into and processed in the pipeline. Stages B, ...pipeline, in accordance with the preferred embodiments of the present invention, to "fill up" empty processing stages is highly advantageous since the processing stages in the pipeline thereby become decouple from one another. In other words, even though... ..data can be transferred into the pipeline and between stages even when one or more processing stages is blocked.

In the embodiment shown in Fig. 1, it is assumed that the...propagate all the way back to the beginning of the pipeline if there is some intermediate stage that is able to accept new data.

In the embodiment illustrated in Fig. 1...has been mentioned. It is to be further understood that each pipeline stage may also process the data it has received arbitrarily before passing it between its internal storage elements or... the portion of the pipeline that contains input and output storage elements and that arbitrarily processes data stored in its storage elements.

Furthermore, the "device" downstream from the ...valid data, but also when a stage requires more than one clock phase to finish processing its data. This also can occur when it creates valid data in one or both... control the passage of data between adjacent storage elements. The VALID signal may also be processed in an analogous manner.

A great advantage of the two-wire interface (one wire for... In addition, two extra latches and a small number of gates are preferably added to process the ACCEPT and VALID signals that are associated with the data latches in each half... application so requires. The interface in accordance with this embodiment can also be used to process analog signals.

As discussed previously, while other conventional timing arrangements may be used, the interface... circuit B1, which may be provided to convert output data from input latch LDIN into intermediate data, which is then later loaded in an output data latch LDOUT, which comprises the... is connected either directly as an input to the validation output latch LVOUT, or via intermediate logic devices or circuits that may alter the signal.

Similarly, the output validation signal QVOUT... to the input of the validation input latch QVIN of the following stage, or via intermediate devices or logic circuits, which may alter the validation signal. This output QVIN is ...word.

Preferred Data Structure - "tokens"

In the sample application shown in Fig. 4, each stage processes all input data, since there is no control circuitry that excludes any stage from allowing... are connected together in a relatively simple configuration. The simplest configuration is a pipeline of processing steps. For example, in the one shown in Fig. 1. The use of tokens, however...flows from left to right in the diagram. Data enters the machine and passes into processing Stage A: This may or may not modify the data and it then passes the... advantage of the tokens is their ability to achieve this kind of communication. Since any processing stage that does not recognize a token simply passes it on unaltered to the next... is transmitted along with the address and data fields in each token so that a processing stage can pass on a token (which can be of arbitrary length) without having to... be the first word of a new token.

Note that although the simple pipeline of processing stages is particularly useful, it will be appreciated that tokens may be applied to more complicated configurations of processing elements. An example of a more complicated processing element is described below.

It is not necessary, in accordance with the present invention, to... has extension bits. An example of this is a token that activates a stage that processes video quantization values stored in a quantization table (typically a memory device). For example, a...turn, is of great importance in video data pipeline systems since it ensures that all processing stages can be continuously running at full bandwidth.

In accordance to the present invention, in... some other chips in the set. This is advantageous both from the perspective of a customer and from that of a chip manufacturer. Even if modifications mean that all chips are...the end of a token (and hence the start of the next token) to be processed correctly (including simple non-manipulative transfer), even if the token is not recognized by the... block diagram of a pipeline stage whose function is as follows. If the stage is processing a predetermined token (known in this example as the DATA token), then it will duplicate... the address field of the DATA token. If, on the other hand, the stage is processing any other kind of token, it will delete every word. The overall effect is that... respective output signals:

In the duplication stage, the output from the data latch LDIN forms intermediate data referred to as MID(underscore)DATA. This intermediate data word is loaded into the data output latch LDOUT only when an intermediate acceptance signal (labeled "MID(underscore)ACCEPT" in Fig. 8a) is set HIGH.

The portion of... ..data. These include a "DATA(underscore)TOKEN" signal that indicates that the circuitry is currently processing a valid DATA Token, and a NOT(underscore)DUPLICATE signal which is used to control duplication of data. When the circuitry is processing a DATA Token, the NOT(underscore)DUPLICATE signal toggles between a HIGH and a LOW... ..the token to be duplicated once (but no more times). When the circuitry is not processing a valid DATA Token then the NOT(underscore)DUPLICATE signal is held in a HIGH state. Accordingly, this means that the token words that are being processed are not duplicated.

As Fig. 8a illustrates, the upper six bits of 8-bit intermediate data word and the output signal QI1 from the latch LI1 form inputs to a... ..explained further below.

Latch LO1 performs the function of latching the last value of the intermediate extension bit (labeled "MID(underscore)EXTN" and as signal S4), and it loads this value... ..and the DATA(underscore)TOKEN signal will become "0", indicating that the circuitry is not processing a DATA token.

If QI1 is "0" and SO is "0", thereby indicating a DATA... ..phase and the DATA(underscore)TOKEN signal will become "1", indicating that the circuitry is processing a DATA token.

The NOT(underscore)DUPLICATE signal (the output signal QO3) is similarly loaded...LVOUT at the same time that MID(underscore)DATA is loaded into LDOUT and the intermediate extension bit (signal S4) is loaded into LEOUT. Signal S5 is also combined with the...above. This has the effect that all tokens except the one that causes the duplication process will be deleted from the token stream, since a device connected to the output terminals (OUTDATA, OUTEXTN and OUTVALID) will not recognize these token words as valid data.

As before... ..and is duplicated.

Referring now more particularly to Figure 10, there is shown a reconfigurable process stage in accordance with one aspect of the present invention.

Input latches 34 receive an... ..the input latches 34 is passed as a first input over line 35 to a processing unit 36. A first output from the token decode subsystem 33 is passed over line 37 as a second input to the processing unit 36. A second output from the token decode 33 is passed over line 40 to an action identification unit 39. The action identification unit 39 also receives input from registers 43 and 44 over line 46. The registers 43... ..is determined by the history of tokens previously received. The output from the action identification unit 39 is passed over line 38 as a third input to the processing unit 36. The output from the processing unit 36 is passed to output latches 41. The output from the output latches 41 is... ..decoder 56 is passed over line 63 as an input to an Index to Data Unit (ITOD) 64. The Huffman decoder 56 and the ITOD 64 work together as a single logical unit. The output from the ITOD 64 is passed over line 65 to an arithmetic logic unit (ALU) 66. A first output from the ALU 66 is passed over line 67 to...blocks 133.

Referring to Figure 14b, in the JPEG and H.261 standards, the Common Intermediate Format (CIF) is used, wherein a picture 141 is encoded as 6 rows each containing... ..in a zigzag direction indicated by the arrow 144. The GOBs 142 are, in turn, processed row-by-row, left-to-right in each row.

Referring now to Figure 14c, it... ..in accordance with the practice of the present invention. A first picture 161 to be processed contains a first PICTURE(underscore)START token 162, first-picture information of indeterminate length 163, and a first PICTURE(underscore)END token 164. A second picture 165 to be processed contains a second PICTURE(underscore)START token 166, second picture information of indeterminate length 167... ..tokens 162

and 166 indicate the start of the pictures 161 and 165 to the processor. Likewise, the PICTURE(underscore)END tokens 164 and 168 signify the end of the pictures 161 and 165 to the processor. This allows the processor to process picture information 163 and 167 of variable lengths.

Referring to Figure 17, a split 171...Video Formatter (not shown in Figure 17).

Referring now to Figure 18, the prediction filtering process is illustrated. A forward picture 201 is passed over line 202 as a first input...the right of the value decode shift register 230, as indicated by area 231. This process eliminates overlapping start code images, as discussed below. A first output from the value decode...Code Detector. The Start Code Detector then receives a first data value image 244. Before processing the first data value image 244, the Start Code Detector may detect a second start...image 244 at a length 246. If this occurs, the Start Code Detector does not process the first data value image 244, and instead receives and processes a second data...line 1 of Table 600, whenever a "sequence start" image is received during H.261 processing or a "picture start" image is received during MPEG processing, the entire group of four control tokens is generated, each followed by its corresponding data...Picture Decoding

3. Motion Picture Decompression

4. RAM Memory Map

5. Bitstream Characteristics

6. Reconfigurable Processing Stage

7. Multi-Standard Coding

8 Multi-Standard Processing Circuit-2nd Mode of Operation

9. Start Code Detector

10. Tokens

11. DRAM Interface

12...described herein in greater detail) and reformatting this output for use, including display in a computer or other display systems, including a video display system. Implementation of this formatting varies significantly...the Spatial Decoder circuits.

The Spatial Decoder of the present invention performs all the required processing within a single picture. This reduces the redundancy within one picture.

The Temporal Decoder reduces...modeller 75, the inverse zig-zag 81 and the inverse DCT 83. The standard independent units within the Huffman decoder and parser include the ALU 66 and the token formatter 71.

Referring now to Figure 12, the standard-independent units include the DRAM interface 100, the fork 91, the FIFO register 96, the summer 98 and the output selector 106. The standard dependent units are the address generator 94, which is different in H.261 and in MPEG, and...much of the operation is very similar between the three different compression standards.

The next unit is the state machine 68 (Figure 11) located within the Huffman decoder and parser. Here...The same holds true for JPEG, which is a third, completely independent program.

The next unit is the Huffman decoder 56 which functions with the index to data unit 64. Those two units cooperate together to perform the Huffman decoding. Here, the algorithm that is used for Huffman...to the Huffman decoder at different times consistent with the standard in operation.

The last unit on the chip that is dependent on the compression standard is the inverse quantizer 79...an H.261 group of blocks and an MPEG slice. When H.261 data is processed after the Start Code Detector, each group of blocks is preceded by a slice(underscore... ..these standards have totally different sets of tables.

As previously indicated, most of the system units are compression standard independent. If a unit is standard independent, and such units need not remember what CODING(underscore)STANDARD is being processed. All of the units that are standard dependent remember the compression standard as the CODING(underscore)STANDARD token flows... ..CODING(underscore)STANDARD tokens at the Start Code Detector that is positioned as the first unit in the pipeline, this change of compression standard is readily handled. The token says a... ..found in the standard, i.e. from the bitstream into a prediction mode token. This processing is performed by the Huffman decoder and parser state machine, where it is easy to... ..to that token. By having these tokens and using them appropriately, the design of other units in the machine is simplified. Although there may be some complications in the program, benefits... ..a first encoded signal (the MPEG or H.261 encoded video signal) in a pipeline processing system. The Temporal Decoder is not needed for JPEG decoding.

In this regard, the invention... ..the use of a single pipeline decoder and decompression system. The decoding and decompression pipeline processor is organized on a unique and special configuration which allows the handling of the multi... ..video signals through the use of techniques all compatible with the single pipeline decoder and processing system. The Spatial Decoder is combined with the Temporal Decoder, and the Video Formatter is...with only still pictures. The compression standard independent Spatial Decoder performs all of the data processing within the boundaries of a single picture. Such a decoder handles the spatial decompression of... ..to the multi-standard, configurable Video Formatter, which then provides an output to the display terminal. In a first sequence of similar pictures, each decompressed picture at the output of the... ..of control tokens and DATA tokens, in combination with a plurality of sequentially-positioned reconfigurable processing stages selected and organized to act as a standard-independent, reconfigurable-pipeline-processor.

With regard to JPEG decoding, a single Spatial Decoder with no off chip DRAM can... ..video. Accordingly, signals carried by DATA tokens pass directly through the Temporal Decoder without further processing when the Temporal Decoder is configured for a JPEG operation.

Another aspect of the present... ..for subsequent use in temporal decoding of subsequent pictures.

Generally, the Temporal Decoder performs the processing between pictures either earlier and/or later in time with reference to the picture currently... ..is distributed among several areas of DRAM in the sense that the decompressed output information, processed by the Spatial Decoder, is stored in other DRAM registers by other random access memories...first decoder circuit (the Spatial Decoder) directly to the Video Formatter for handling without signal processing delay.

The Temporal Decoder also reorders the blocks of picture data for display by a... ..from a selection of pictures which have arrived earlier or later than the picture under processing. When a picture is described in this context, it may mean any one of the... ..2. The result, i.e., the final decoded picture resulting from the addition of a process step performed by the decoder;

3. Previously decoded pictures read from the DRAM; and

4... ..START token and a subsequent PICTURE(underscore)END token.

After the picture data information is processed by the Temporal Decoder, it is either displayed or written back into a picture memory location. This information is then kept for further reference to be used in processing another different coded data picture.

Re-ordering of the MPEG encoded pictures for visual display...used to encode a referenced picture of a picture might be identified as being one unit long, another picture might be a number of units long, while still a third picture could be a fraction of that unit.

None of the existing standards (MPEG 1.2, JPEG, H.261) define a way of... picture rate, whereas the Video Formatter can handle a variable input picture rate.

6. RECONFIGURABLE PROCESSING STAGE

Referring again to Figure 10, the reconfigurable processing stage (RPS) comprises a token decode circuit 33 which is employed to receive the tokens... input latches 34. The output of the token decode circuit 33 is applied to a processing unit 36 over the two-wire interface 37 and an action identification circuit 39. The processing unit 36 is suitable for processing data under the control of the action identification circuit 39. After the processing is completed, the processing unit 36 connects such completed signals to the output, two-wire interface bus 40 through output... token decode circuit 33 are applied simultaneously to the action identification circuit 39 and the processing unit 36. The action identification function as well as the RPS is described in further detail... not standard independent circuits. The data flows through the token decode circuit 33, through the

processing unit 36 and onto the two-wire interface circuit 42 through the output latches 41. If... wire interface 42 through the output circuit 41. The present invention operates as a pipeline processor having a two-wire interface for controlling the movement of control tokens through the pipeline... time, the token decode circuit 33 provides a proper flag or index signal to the processing unit 36 to alert it to the presence of the token being handled by the action identification circuit 39.

Control tokens may also be processed.

A more detailed description of the various types of tokens usable in the present invention... standard now passing through the state machine shown with reference to Figure 10.

Similarly, the processing unit 36 which is under the control of the action identification circuit 39 is now ready to process the information contained in the data fields of the DATA token when it is appropriate... action identification circuit 39 and is immediately followed by a DATA token which is then processed by the processing unit 36. The control token exits the output latches circuit 41 over the output two-wire interface 42 immediately preceding the DATA token which has been processed within the processing unit 36.

In the present invention, the action identification circuit, 39, is a state machine holding... show that the action can also be affected by the token that is currently being processed by the token decode circuit 33.

In general, there is shown token decoding and data processing in accordance with the present invention. The data processing is performed as configured by the action identification circuit 39. The action is affected by... information stored from previously decoded tokens in registers 43 and 44, the current token under processing, and the state and history information that the action identification unit 39 has itself acquired. A distinction is thereby shown between Control tokens and DATA tokens.

In any RPS, some tokens are viewed by that RPS unit as being Control tokens in that they affect the operation of the RPS presumably at... are viewed by the RPS as DATA tokens. Such DATA tokens contain information which is processed by the RPS in a way that is determined by the design of the particular... view of the same token. Some of the tokens might be viewed by one RPS unit as DATA Tokens while another RPS unit might decide that it is actually a Control Token. For example, the quantization table information... into a token called a quantization table token (QUANT(underscore)TABLE) which goes down the processing pipeline. As far as that machine is concerned, all of that was data; it was... sort of data into another sort of data, which is clearly a function of the processing performed by that portion of the machine. However, when that information gets to the inverse... present. This

information is viewed as control information, and then that control information affects the processing that is done on subsequent DATA tokens because it affects the number that you multiply... ..important feature of the invention is that each of the stages of circuitry has the processing capability within it to be able to perform the necessary operations for each of the... ..operations are to be performed at a given time, come as tokens. There is one processing element that differs between the different stages to provide this capability. In the state machine...standard is and it looks up the parameters that it needs to apply to the processing elements in order to perform a proper operation. For example, the inverse quantizer will look... ..is set to 1 for a particular compression standard, and will apply that to its processing circuitry.

In a similar sense the Huffman decoder 56 has a number of tables within... ..MPEG video standard or the JPEG video standard. These three compression coding standards specify similar processes to be done on the arriving data, but the structure of the datastreams is different... ..token stream embodying the current coding standard. The control tokens are passed through the pipeline processor, and are used. i.e., decoded, in the state machines to which they are relevant... ..this regard, the DATA Tokens are treated in the same fashion, insofar as they are processed only in the state machines that are configurable by the control tokens into processing such DATA Tokens. In the remaining state machines, they pass through unchanged.

More specifically, a... ..signals. The remaining portions of the token are used to indicate and identify the internal processing control function which is standard for all of the datastreams passing through the pipeline processor. In one form of the invention, the token extension is used to carry the current...accompanying data. As previously discussed, this information is utilized in the system to reconfigure the processing stage used to perform the function required by the various standards created for that purpose... ..picture number as indicated by the value.

The system also includes a multi-stage parallel processing pipeline operating under the principles of the two-wire interface previously described. Each of the... ..the token presently entering the state machine into the action identification circuit 39 or the processing unit 36, as appropriate. The processing unit has been previously reconfigured by the next previous control token into the form needed for handling the current coding standard, which is now entering the processing stage and carried by the next DATA token. Further, in accordance with this aspect of the invention, the succeeding state machines in the processing pipeline can be functioning under one coding standard, i.e., H.261, while a previous stage... ..tokens required to decode a number of coding standards with a fixed number of reconfigurable processing stages. More specifically, the PICTURE(underscore)END control token is employed because it is important... ..standard machine, it is necessary to create additional control tokens within the multi-standard pipeline processing machine which will then indicate which one of the standard decoding techniques to use. Such... ..and to push the current picture through the decoder to the display.

8. MULTI-STANDARD PROCESSING CIRCUIT - SECOND MODE OF OPERATION

A compression standard-dependent circuit, in the form of the...of the Start Code Detector will subsequently be discussed in further detail, as will the process of starting up of the decoder.

The aforementioned description ...the data which immediately follows according to the standard. However, in the multi-standard pipeline processing system of the present invention, where compatibility is required for multiple standards, the system has... ..signals, including flag signals, are generated by each state machine to handle some of the processing within that state machine. Values carried in the standards can be used to access machine... ..its contents must be removed from the two wire interface to ensure that no further processing takes place using these 3 bytes. The decode register is emptied, and the value decode... ..10. TOKENS

In the practice of the present invention, a token is a universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions and is adapted for use with a reconfigurable processing stage (RPS) which is a stage, which in response to a recognized token, reconfigures itself to perform various operations.

Tokens may be either position dependent or position independent upon the processing stages for performance of various functions. Tokens may also be metamorphic in that they can be altered by a processing stage and then passed down the pipeline for performance of further functions. Tokens may interact... other functions, and the specific interaction with a stage may be conditioned by the previous processing history of a stage.

A PICTURE(underscore)END token is a way of signalling the... through a fixed size, fixed width buffer.

The present invention is directed to a pipeline processing system which has a variable configuration which uses tokens and a two-wire system. The... do not use control tokens.

The control tokens are generated by circuitry within the decoder processor and emulate the operation of a number of different type standard-dependent signals passing into the serial pipeline processor for handling. The technique used is to study all the parameters of the multi-standards that are selected for processing by the serial processor and noting 1) their similarities, 2) their dissimilarities, 3) their needs and requirements and 4) selecting the correct token function to effectively process all of the standard signals sent into the serial processor. The functions of the tokens are to emulate the standards. A control token function is... the standard dependent signals and as an element to transmit control information through the pipeline processor.

In prior art system, a dedicated machine is designed according to well-known techniques to... tokens provide and make a sensible format for communicating information through the decompression circuit pipeline processor. In the design selected hereinafter and used in the preferred embodiment, each word of a... However, this is not a limitation on the invention, but on the magnitude of the processing steps elected to be accomplished by use of these tokens. It is to be noted... bit address for use in accessing the random access memories used throughout this serial decompression processor. This provides an additional degree of variability that facilitates a broad range of versatility.

As previously described, the DATA token carries data from one processing stage to the next. Consequently, the characteristics of this token change as it passes through... longest number of data bits because it needs to provide the most information to the processing

unit so that it can start the decompression with as much information as possible. Words which... to receive an address, it waits for the address generator to supply a valid address, processes that address and then sets the accept line high for one clock period. Thus, it... be read. This signal passes between two asynchronous clock regimes and, therefore, passes through three synchronizing flip flops.

Provided RAM2 312 is empty, the next item of data to arrive on... interesting.

In general, prediction data will be offset from the position of the block being processed as specified in the motion vectors in x and y. Thus, the block of data... address, 9. Data is read from this address and the x value is incremented. The process is repeated until the x value reaches its stop value, at which point, the y... is read, the x value is again incremented until it reaches its stop value. The process is repeated until both x and y values have reached their stop values. Thus, the... invention, is that additional information must be provided to the prediction filters to indicate what processing is required on the data. This consists of the following:

a "last byte" signal indicating... bit 0) is incremented and the x address (3 LSBS) is reset to zero. This process is repeated until 64 bytes have been read. With a 16 or 32 bit wide... register while its access register is set to zero, the results are undefined.

14. MICRO-PROCESSOR INTERFACE

A standard byte wide micro-processor interface (MPI) is used on all circuits with in the Spatial Decoder and Temporal Decoder... ..the parameter column. The actual specifications are shown in the respective columns min, max and units.

The DC operating conditions can be seen with reference to Table A.6.3. Here... ..the signal is present the maximum amount of time that this signal is available. The Units column gives the units of measurement used to describe the signals.

16. MPI WRITE TIMING

The general description of...Consequently, the machine will not go into error recovery mode and will successfully continue to process the coded data.

A still further advantage of the use of a PICTURE(underscore)END token is that the serial pipeline processor will continue the processing of uninterrupted data. Through the use of a PICTURE(underscore)END token, the serial pipeline processor is configured to handle less than the expected amount of data and, therefore, continues processing. Typically, a prior art machine would stop itself because of an error condition. As previously... ..of the Huffman decode and Video Demultiplexor know the number of blocks that it will process during each picture recovery cycle. When the correct number of blocks do not arrive from... ..Each of the state machines recognizes a FLUSH control token as information not to be processed. Accordingly, the FLUSH token is used to fill up all of the remaining empty parts is like padding for buffers.

The Token Decoder in the Huffman circuit recognizes the FLUSH token and ignores the pseudo... ..less information than normally expected to decode the last picture. The Huffman decode circuit finishes processing the information contained in the last picture, and outputs this information through the DRAM interface... ..token, in accordance with the present invention, is used to pass through the entire pipeline processor and to ensure that the buffers are emptied and that other circuits are reconfigured to... ..underscore)END token, a padding word and a FLUSH token indicating to the serial pipeline processor that the picture processing for the current picture form is completed. Thereafter, the various state machines need reconfiguring to... ..FLUSH token resets each stage as it passes through, but-allows subsequent stages to continue processing. This prevents a loss of data. In other words, the FLUSH token is a variable... ..AFTER PICTURE

The STOP(underscore)AFTER(underscore)PICTURE function is employed to shut down the processing of the serial pipeline decompressing circuit at a logical point in its operation. At this... ..a picture, the STOP(underscore)AFTER(underscore)PICTURE operation signals the end of all current processing.

22. MULTI(underscore)STANDARD - SEARCH MODE

Another feature of the present invention is the use... ..underscore)MODE control token which is used to reconfigure the input to the serial pipeline processor to look at the incoming bit stream. When the search mode is set, the Start... ..combination of control tokens, and DATA tokens along with the reconfiguration circuits, to provide similar processing.

The use of search mode in the present invention is convenient in many situations including... ..video disc. In general, a search mode is convenient when the user interrupts the normal processing of the serial pipeline at a point where the machine does not expect such an...be the case.

In brief, the Huffman Decoder 321 works in conjunction with the other units shown in Figure 27. These other units are the Parser State Machine 322, the inshifter 323, the Index to Data unit 324, the ALU 325, and the Token Formatter 326. As described previously, connection between these blocks is governed by a two wire interface. A more detailed description of how these units function is subsequently described herein in greater detail, the focus here is on particular aspects... ..control certain functions of the Index to Data 324 and ALU 325. Control of these

units by the Huffman Decoder is necessary for proper decoding of block-level information. Having the... further detail in the "More Detailed Description of the Invention" section.

The Index to Data unit 324 performs the second part of the multi-part algorithm. This unit contains a look up table that provides the actual Huffman decoded data. Entries in the...by detecting these in the Huffman Decoder 321, rather than in the Index to Data unit 324.

This index number is then passed to the Index to Data unit 324. In essence, the Index to Data unit is a look-up table. In accordance with one aspect of the algorithm, the look... format that JPEG specifies for transferring an alternate JPEG table.

From the Index to Data unit 324, the decoded index number or other data is passed, together with the accompanying control... the entering data to ensure that the DATA tokens are of the correct size for processing. In fact, the token stream can be corrected in some situations if the error is... an order that is useful for the decompression circuits, but not for the particular display unit being used. When a block of data enters the Buffer Manager, the Buffer Manager supplies... the output of the Spatial Decoder or Temporal Decoder and re-format it for a computer or display system. The details of this formatting will vary between applications. In a simple... Token. The DATA Token can have as many bits as are necessary for carrying out processing at a particular place in the system. All other Tokens ignore the extra bits.

A.3.2 The DATA Token

The DATA Token carries data from one processing stage to the next. Consequently, the characteristics of this Token change as it passes through... will be sufficient to collect DATA Tokens and to detect a few Tokens that provide synchronization information (such as PICTURE(underscore)START). In this regard, see subsequent sections A.16, "Connecting... from the data stream. This provides an alternative to doing the configuration via the micro processor interface.

A.3.4 Description of Tokens

This section documents the Tokens which are implemented... 3.5.1. Note: JPEG requires a 2:1:1 structure for its macroblocks when processing 4:2:2 data. See Table A.3.5.

A.3.6 Special Token formats... either is low then the interface is taken to high impedance.

Note: on-chip data processing is not terminated when the DRAM interface is at high impedance. Therefore, errors will occur... decoded video's picture rate. Accordingly, this clock can be used to provide audio/video synchronization.

A.7.1 Spatial Decoder clock signals

The Spatial Decoder has two different (and potentially... in accordance with the present invention must know what video standard is being input for processing. Thereafter the system can accept either pre-existing Tokens or raw byte data which is... time a value is written into coded(underscore)data (7:0). Software is responsible for settling coded(underscore)extn to 0 before the last word of any Token is written to... 0). The start of this new DATA Token then passes into the Spatial Decoder for processing.

Each time a new 8 bit value is written to coded(underscore)data (7:0... Detector analyses data in the DATA Tokens bit serially. The Detector's normal rate of processing is one bit per clock cycle (of coded(underscore)clock). Accordingly, it will typically decode a byte of coded data every 8 cycles of coded(underscore)clock. However, extra processing cycles are occasionally required, e.g., when a non-DATA Token is supplied or when this clock can be asynchronous to the main decoder(underscore)clock. Data transfer is synchronized to decoder(underscore)clock on-chip.

SECTION A.11 Start code detector

A.11.1... ..Code Detector. So, accessing these registers will be unreliable if the Start Code Detector is processing data. The user is responsible for ensuring that the Start Code Detector is halted before... ..Detector. In this case, the Tokens are passed through the Start Code Detector with no processing to other stages of the Spatial Decoder. These Tokens can only be inserted just before...result will be unpredictable if this is done when the Start Code Detector is actively processing data.

Discard all mode can be safely initiated after any of the Start Code Detector...the start code non-alignment interrupt is suppressed.

In contrast, however, JPEG was designed for a computer environment where byte alignment is guaranteed. Therefore, marker codes should only be detected when byte... ..the other hand, was designed to meet the needs of both communications (bit serial) and

computer (byte oriented) systems. Start codes in MPEG data should normally be byte aligned. However, the...result will be unpredictable if this is done when the Start Code Detector is actively processing data. So, before initiating a start code search, the Start Code Detector should be stopped so no data is being processed. The Start Code Detector is always in this condition if any of the Start Code...the spatial video decoding circuits (inverse modeler, quantizer and DCT). This second logical buffer allows processing time to include a spread so as to accommodate processing pictures having varying amounts of data.

Both buffers are physically held in a single off in A.13.1.1, the unit for all the above mentioned registers is a 512 bit block of data. Accordingly, the... ..until there is space in the buffer. If a buffer continues to be full, more processing stages "up stream" of the buffer will halt until the Spatial Decoder is unable to... ..converting coded data into Tokens started by the Start Code Detector. There are four main processing blocks in the Video Demux: Parser State Machine, Huffman decoder (including an ITOD), Macroblock counter... ..or state machine follows the syntax of the coded video data and instructs the other units. The Huffman decoder converts variable length coded (VLC) data into integers. The Macroblock counter keeps...

Specification: ...pipeline, in accordance with the preferred embodiments of the present invention, to "fill up" empty processing stages is highly advantageous since the processing stages in the pipeline thereby become decouple from one another. In other words, even though... ..data can be transferred into the pipeline and between stages even when one or more processing stages is blocked.

In the embodiment shown in Fig. 1, it is assumed that the...propagate all the way back to the beginning of the pipeline if there is some intermediate stage that is able to accept new data.

In the embodiment illustrated in Fig. 1...has been mentioned. It is to be further understood that each pipeline stage may also process the data it has received arbitrarily before passing it between its internal storage elements or... ..the portion of the pipeline that contains input and output storage elements and that arbitrarily processes data stored in its storage elements.

Furthermore, the "device" downstream from the pipeline Stage F... ..valid data, but also when a stage requires more than one clock phase to finish processing its data. This also can occur when it creates valid data in one or both... ..control the passage of data between adjacent storage elements. The VALID signal may also be processed in an analogous manner.

A great advantage of the two-wire interface (one wire for... ..In addition, two extra latches and a small number of gates are preferably added to process the ACCEPT and VALID signals that are associated with the data latches in each half... ..application so requires. The interface in accordance with this embodiment can also be used to process analog signals.

As discussed previously, while other conventional timing arrangements may be used, the interface... circuit B1, which may be provided to convert output data from input latch LDIN into intermediate data, which is then later loaded in an output data latch LDOUT, which comprises the as an input to the validation output latch LVOUT, or via intermediate logic devices or circuits that may alter the signal.

Similarly, the output validation signal QVOUT... to the input of the validation input latch QVIN of the following stage, or via intermediate devices or logic circuits, which may alter the validation signal. This output QVIN is also...word.

Preferred Data Structure - "tokens"

In the sample application shown in Fig. 4, each stage processes all input data, since there is no control circuitry that excludes any stage from allowing...are connected together in a relatively simple configuration. The simplest configuration is a pipeline of processing steps. For example, in the one shown in Fig. 1. The use of tokens, however... flows from left to right in the diagram. Data enters the machine and passes into processing Stage A. This may or may not modify the data and it then passes the... advantage of the tokens is their ability to achieve this kind of communication. Since any processing stage that does not recognize a token simply passes it on unaltered to the next... is transmitted along with the address and data fields in each token so that a processing stage can pass on a token (which can be of arbitrary length) without having to... be the first word of a new token.

Note that although the simple pipeline of processing stages is particularly useful, it will be appreciated that tokens may be applied to more complicated configurations of processing elements. An example of a more complicated processing element is described below.

It is not necessary, in accordance with the present invention, to... has extension bits. An example of this is a token that activates a stage that processes video quantization values stored in a quantization table (typically a memory device). For example, a...turn, is of great importance in video data pipeline systems since it ensures that all processing stages can be continuously running at full bandwidth.

In accordance to the present invention, in... some other chips in the set. This is advantageous both from the perspective of a customer and from that of a chip manufacturer. Even if modifications mean that all chips are...the end of a token (and hence the start of the next token) to be processed correctly (including simple non-manipulative transfer), even if the token is not recognized by the... block diagram of a pipeline stage whose function is as follows. If the stage is processing a predetermined token (known in this example as the DATA token), then it will duplicate... the address field of the DATA token. If, on the other hand, the stage is processing any other kind of token, it will delete every word. The overall effect is that... signal IN VALID.

In the duplication stage, the output from the data latch LDIN forms intermediate data referred to as MID(underscore)DATA. This intermediate data word is loaded into the data output latch LDOUT only when an intermediate acceptance signal (labeled "MID(underscore)ACCEPT" in Fig. 8a) is set HIGH.

The portion of... data. These include a "DATA(underscore)TOKEN" signal that indicates that the circuitry is currently processing a valid DATA Token, and a NOT(underscore)DUPLICATE signal which is used to control duplication of data. When the circuitry is processing a DATA Token, the NOT(underscore)DUPLICATE signal toggles between a HIGH ...the token to be duplicated once (but no more times). When the circuitry is not processing a valid DATA Token then the NOT(underscore)DUPLICATE signal is held in a HIGH state. Accordingly, this means that the token words that are being processed are not duplicated.

As Fig. 8a illustrates, the upper six bits of 8-bit intermediate data word and the output signal QI1 from the latch LI1 form inputs to a... is explained further below.

Latch LO1 performs the function of latching the last value of the intermediate extension bit (labeled "MID(underscore)EXTN" and as signal S4), and it loads this value... ..and the DATA(underscore)TOKEN signal will become "0", indicating that the circuitry is not processing a DATA token.

If QI1 is "0" and SO is "0", thereby indicating a DATA... ..phase and the DATA(underscore)TOKEN signal will become "1", indicating that the circuitry is processing a DATA token.

The NOT(underscore)DUPLICATE signal (the output signal QO3) is similarly loaded...LVOUT at the same time that MID(underscore)DATA is loaded into LDOUT and the intermediate extension bit (signal S4) is loaded into LEOUT. Signal S5 is also combined with the...above. This has the effect that all tokens except the one that causes the duplication process will be deleted from the token stream, since a device connected to the output terminals (OUTDATA, OUTEXTN and OUTVALID) will not recognize these token words as valid data.

As before... ..and is duplicated.

Referring now more particularly to Figure 10, there is shown a reconfigurable process stage.

Input latches 34 receive an input over a first bus 31. A first output from... ..the input latches 34 is passed as a first input over line 35 to a processing unit 36. A first output from the token decode subsystem 33 is passed over line 37 as a second input to the processing unit 36. A second output from the token decode 33 is passed over line 40 to an action identification unit 39. The action identification unit 39 also receives input from registers 43 and 44 over line 46. The registers 43... ..is determined by the history of tokens previously received. The output from the action identification unit 39 is passed over line 38 as a third input to the processing unit 36. The output from the processing unit 36 is passed to output latches 41. The output from the output latches 41 is... ..decoder 56 is passed over line 63 as an input to an Index to Data Unit (ITOD) 64. The Huffman decoder 56 and the ITOD 64 work together as a single logical unit. The output from the ITOD 64 is passed over line 65 to an arithmetic logic unit (ALU) 66. A first output from the ALU 66 is passed over line 67 to...blocks 133.

Referring to Figure 14b, in the JPEG and H.261 standards, the Common Intermediate Format (CIF) is used, wherein a picture 141 is encoded as 6 rows each containing... ..in a zigzag direction indicated by the arrow 144. The GOBS 142 are, in turn, processed row-by-row, left-to-right in each row.

Referring now to Figure 14c, it... ..in accordance with the practice of the present invention. A first picture 161 to be processed contains a first PICTURE(underscore)START token 162, first-picture information of indeterminate length 163, and a first PIC-TURE(underscore)END token 164. A second picture 165 to be

processed contains a second PICTURE(underscore)START token 166, second picture information of indeterminate length 167... ..tokens 162 and 166 indicate the start of the pictures 161 and 165 to the processor. Likewise, the PICTURE(underscore)END tokens 164 and 168 signify the end of the pictures 161 and 165 to the processor. This allows the processor to process picture information 163 and 167 of variable lengths.

Referring to Figure 17, a split 171...Video Formatter (not shown in Figure 17).

Referring now to Figure 18, the prediction filtering process is illustrated. A forward picture 201 is passed over line 202 as a first input... ..the right of the value decode shift register 230, as indicated by area 231. This process eliminates overlapping start code images, as discussed below. A first output from the value decode... ..Code Detector. The Start Code Detector then receives a first data value image 244. Before processing the first data value image 244, the Start Code Detector may detect a second start...image 244 at a length 246. If this occurs, the Start Code Detector does not process the first data value image 244, and instead receives and processes a second data value image 247.

Referring now to Figure 22, a flag generator 251... line 1 of Table 600, whenever a "sequence start" image is received during H.261 processing or a "picture start" image is received during MPEG processing, the entire group of four control tokens is generated, each followed by its corresponding data...Picture Decoding

3. Motion Picture Decompression

4. RAM Memory Map

5. Bitstream Characteristics

6. Reconfigurable Processing Stage

7. Multi-Standard Coding

8. Multi-Standard Processing Circuit-2nd Mode of Operation

9. Start Code Detector

10. Tokens

11. DRAM Interface

12... described herein in greater detail) and reformatting this output for use, including display in a computer or other display systems, including a video display system. Implementation of this formatting varies significantly... the Spatial Decoder circuits.

The Spatial Decoder of the present invention performs all the required processing within a single picture. This reduces the redundancy within one picture.

The Temporal Decoder reduces... modeller 75, the inverse zig-zag 81 and the inverse DCT 83. The standard independent units within the Huffman decoder and parser include the ALU 66 and the token formatter 71.

Referring now to Figure 12, the standard-independent units include the DRAM interface 100, the fork 91, the FIFO register 96, the summer 98 and the output selector 106. The standard dependent units are the address generator 94, which is different in H.261 and in MPEG, and...much of the operation is very similar between the three different compression standards.

The next unit is the state machine 68 (Figure 11) located within the Huffman decoder and parser. Here... The same holds true for JPEG, which is a third, completely independent program.

The next unit is the Huffman decoder 56 which functions with the index to data unit 64. Those two units cooperate together to perform the Huffman decoding. Here, the algorithm that is used for Huffman... to the Huffman decoder at different times consistent with the standard in operation.

The last unit on the chip that is dependent on the compression standard is the inverse quantizer 79...an H.261 group of blocks and an MPEG slice. When H.261 data is processed after the Start Code Detector, each group of blocks is preceded by a slice(underscore... these standards have totally different sets of tables.

As previously indicated, most of the system units are compression standard independent. If a unit is standard independent, and such units need not remember what CODING(underscore)STANDARD is being processed. All of the units that are standard dependent remember the compression standard as the CODING(underscore)STANDARD token flows... CODING(underscore)STANDARD tokens at the Start Code Detector that is positioned as the first unit in the pipeline, this change of compression standard is readily handled. The token says a... found in the standard, i.e. from the bitstream into a prediction mode token. This processing is performed by the Huffman decoder and parser state machine, where it is easy to... to that token. By having these tokens and using them appropriately, the design of other units in the machine is simplified. Although there may be some complications in the program,

benefits... ...a first encoded signal (the MPEG or H.261 encoded video signal) in a pipeline processing system. The Temporal Decoder is not needed for JPEG decoding.

In this regard, the invention... ...the use of a single pipeline decoder and decompression system. The decoding and decompression pipeline processor is organized on a unique and special configuration which allows the handling of the multi... ...video signals through the use of techniques all compatible with the single pipeline decoder and processing system. The Spatial Decoder is combined with the Temporal Decoder, and the Video Formatter is...with only still pictures. The compression standard independent Spatial Decoder performs all of the data processing within the boundaries of a single picture. Such a decoder handles the spatial decompression of... ...to the multi-standard, configurable Video Formatter, which then provides an output to the display terminal. In a first sequence of similar pictures, each decompressed picture at the output of the... ...of control tokens and DATA tokens, in combination with a plurality of sequentially-positioned reconfigurable processing stages selected and organized to act as a standard-independent, reconfigurable-pipeline-processor.

With regard to JPEG decoding, a single Spatial Decoder with no off chip DRAM can... ...video. Accordingly, signals carried by DATA tokens pass directly through the Temporal Decoder without further processing when the Temporal Decoder is configured for a JPEG operation.

Another aspect of the present... ...for subsequent use in temporal decoding of subsequent pictures.

Generally, the Temporal Decoder performs the processing between pictures either earlier and/or later in time with reference to the picture currently... ...is distributed among several areas of DRAM in the sense that the decompressed output information, processed by the Spatial Decoder, is stored in other DRAM registers by other random access memories...first decoder circuit (the Spatial Decoder) directly to the Video Formatter for handling without signal processing delay.

The Temporal Decoder also reorders the blocks of picture data for display by a... ...from a selection of pictures which have arrived earlier or later than the picture under processing. When a picture is described in this context, it may mean any one of the... ...2. The result, i.e., the final decoded picture resulting from the addition of a process step performed by the decoder;

3. Previously decoded pictures read from the DRAM; and

4... ...START token and a subsequent PICTURE(underscore)END token.

After the picture data information is processed by the Temporal Decoder, it is either displayed or written back into a picture memory location. This information is then kept for further reference to be used in processing another different coded data picture.

Re-ordering of the MPEG encoded pictures for visual display...used to encode a referenced picture of a picture might be identified as being one unit long, another picture might be a number of units long, while still a third picture could be a fraction of that unit.

None of the existing standards (MPEG 1.2, JPEG, H.261) define away of ending... ...picture rate, whereas the Video Formatter can handle a variable input picture rate.

6. RECONFIGURABLE PROCESSING STAGE

Referring again to Figure 10, the reconfigurable processing stage (RPS) comprises a token decode circuit 33 which is employed to receive the tokens... ...input latches 34. The output of the token decode circuit 33 is applied to a processing unit 36 over the two-wire interface 37 and an action identification circuit 39. The processing unit 36 is suitable for processing data under the control of the action identification circuit 39. After the processing is completed, the processing unit 36 connects such completed signals to the output, two-wire interface bus 40 through output... ...token decode circuit 33 are applied simultaneously to the action identification circuit 39 and the

processing unit 36. The action identification function as well as the RPS Is described in further detail... ..not standard independent circuits. The data flows through the token decode circuit 33, through the processing unit 36 and onto the two-wire interface circuit 42 through the output latches 41. If... ..wire interface 42 through the output circuit 41. The present invention operates as a pipeline processor having a two-wire interface for controlling the movement of control tokens through the pipeline... ..time, the token decode circuit 33 provides a proper flag or index signal to the processing unit 36 to alert it to the presence of the token being handled by the action identification circuit 39.

Control tokens may also be processed.

A more detailed description of the various types of tokens usable in the present invention...standard now passing through the state machine shown with reference to Figure 10.

Similarly, the processing unit 36 which is under the control of the action identification circuit 39 is now ready to process the information contained in the data fields of the DATA token when it is appropriate... ..action identification circuit 39 and is immediately followed by a DATA token which is then processed by the processing unit 36. The control token exits the output latches circuit 41 over the output two-wire interface 42 immediately preceding the DATA token which has been processed within the processing unit 36.

In the present invention, the action identification circuit, 39, is a state machine holding... ..show that the action can also be affected by the token that is currently being

processed by the token decode circuit 33.

In general, there is shown token decoding and data processing in accordance with the present invention. The data processing is performed as configured by the action identification circuit 39. The action is affected by... ..information stored from previously decoded tokens in registers 43 and 44, the current token under processing, and the state and history information that the action identification unit 39 has itself acquired. A distinction is thereby shown between Control tokens and DATA tokens.

In any RPS, some tokens are viewed by that RPS unit as being Control tokens in that they affect the operation of the RPS presumably at... ..are viewed by the RPS as DATA tokens. Such DATA tokens contain information which is processed by the RPS in a way that is determined by the design of the particular... ..view of the same token. Some of the tokens might be viewed by one RPS unit as DATA Tokens while another RPS unit might decide that it is actually a Control Token. For example, the quantization table information... ..into a token called a quantization table token (QUANT(underscore)TABLE) which goes down the processing pipeline. As far as that machine is concerned, all of that was data; it was... ..sort of data into another sort of data, which is clearly a function of the processing performed by that portion of the machine. However, when that information gets to the inverse... ..present This information is viewed as control information, and then that control information affects the processing that is done on subsequent DATA tokens because it affects the number that you multiply... ..important feature of the invention is that each of the stages of circuitry has the processing capability within it to be able to perform the necessary operations for each of the... ..operations are to be performed at a given time, come as tokens. There is one processing element that differs between the different stages to provide this capability. In the state machine...standard is and it looks up the parameters that it needs to apply to the processing elements in order to perform a proper operation. For example, the inverse quantizer will look... ..is set to 1 for a particular compression standard, and will apply that to its processing circuitry.

In a similar sense the Huffman decoder 56 has a number of tables within... ..MPEG video standard or the JPEG video standard. These three compression coding standards specify similar processes to be done on the arriving data, but the structure of the datastreams is different... ..token stream embodying the current coding standard. The control tokens are passed through the pipeline processor, and are used, i.e., decoded, in the state machines to which they are

relevant... ..this regard, the DATA Tokens are treated in the same fashion, insofar as they are processed only in the state machines that are configurable by the control tokens into processing such DATA Tokens. In the remaining state machines, they pass through unchanged.

More specifically, a... ..signals. The remaining portions of the token are used to indicate and identify the Internal processing control function which is standard for all of the datastreams passing through the pipeline processor. In one form of the invention, the token extension is used to carry the current... ..accompanying data. As previously discussed, this information is utilized in the system to reconfigure the processing stage used to perform the function required by the various standards created for that purpose... ..picture number as indicated by the value.

The system also includes a multi-stage parallel processing pipeline operating under the principles of the two-wire interface previously described. Each of the... ..the token presently entering the state machine into the action identification circuit 39 or the processing unit 36, as appropriate. The processing unit has been previously reconfigured by the next previous control token into the form needed for handling the current coding standard, which is now entering the processing stage and carried by the next DATA token. Further, in accordance with this aspect of the invention, the succeeding state machines in the processing pipeline can be functioning under one coding standard, i.e., H.261, while a previous... ..tokens required to decode a number of coding standards with a fixed number of reconfigurable processing stages. More specifically, the PICTURE(underscore)END control token is employed because it is important... ..standard machine, it is necessary to create additional control tokens within the multi-standard pipeline processing machine which will then indicate which one of the standard decoding techniques to use. Such... ..and to push the current picture through the decoder to the display.

8. MULTI-STANDARD PROCESSING CIRCUIT - SECOND MODE OF OPERATION

A compression standard-dependent circuit, in the form of the... ..of the Start Code Detector will subsequently be discussed in further detail, as will the process of starting up of the decoder.

The aforementioned description has been concerned primarily with ...the data which immediately follows according to the standard. However, in the multi-standard pipeline processing system of the present invention, where compatibility is required for multiple standards, the system has... ..signals, including flag signals, are generated by each state machine to handle some of the processing within that state machine. Values carried in the standards can be used to access machine... ..its contents must be removed from the two wire interface to ensure that no further processing takes place using these 3 bytes. The decode register is emptied, and the value decode... ..10. TOKENS

In the practice of the present invention, a token is a universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions and is adapted for use with a reconfigurable processing stage (RPS) which is a stage, which in response to a recognized token, reconfigures itself to perform various operations.

Tokens may be either position dependent or position independent upon the processing stages for performance of various functions. Tokens may also be metamorphic in that they can be altered by a processing stage and then passed down the pipeline for performance of further functions. Tokens may interact... ..other functions, and the specific interaction with a stage may be conditioned by the previous processing history of a stage.

A PICTURE(underscore)END token is a way of signalling the... ..through a fixed size, fixed width buffer.

The present invention is directed to a pipeline processing system which has a variable configuration which uses tokens and a two-wire system. The... ..do not use control tokens.

The control tokens are generated by circuitry within the decoder processor and emulate the operation of a number of different type standard-dependent signals passing into the serial pipeline processor for handling. The technique used is to study all the parameters of the multi-standards that are selected for processing by the serial processor and noting 1) their similarities, 2) their dissimilarities, 3) their needs and requirements and 4) selecting the correct token function to effectively process all of the standard signals sent into the serial processor. The functions of the tokens

are to emulate the standards. A control token function is... ..the standard dependent signals and as an element to transmit control information through the pipeline processor.

In prior art system, a dedicated machine is designed according to well-known techniques to... ..tokens provide and make a sensible format for communicating information through the decompression circuit pipeline processor. In the design selected hereinafter and used in the preferred embodiment, each word of a... ..However, this is not a limitation on the invention, but on the magnitude of the processing steps elected to be accomplished by use of these tokens. It is to be noted... ..bit address for use in accessing the random access memories used throughout this serial decompression processor. This provides an additional degree of variability that facilitates a broad range of versatility.

As previously described, the DATA token carries data from one processing stage to the next. Consequently, the characteristics of this token change as it passes through... ..longest number of data bits because it needs to provide the most information to the processing unit so that it can start the decompression with as much information as possible. Words which...to receive an address, it waits for the address generator to supply a valid address, processes that address and then sets the accept line high for one clock period. Thus, it... ..be read. This signal passes between two asynchronous clock regimes and, therefore, passes through three synchronizing flip flops.

Provided RAM2 312 is empty, the next item of data to arrive on...interesting.

In general, prediction data will be offset from the position of the block being processed as specified in the motion vectors in x and y. Thus, the block of data... ..address, 9. Data is read from this address and the x value is incremented. The process is repeated until the x value reaches its stop value, at which point, the y... ..is read, the x value is again incremented until it reaches its stop value. The process is repeated until both x and y values have reached their stop values. Thus, the...invention, is that additional information must be provided to the prediction filters to indicate what processing is required on the data. This consists of the following:

a "last byte" signal indicating... ..bit 0) is incremented and the x address (3 LSBS) is reset to zero. This process is repeated until 64 bytes have been read. With a 16 or 32 bit wide...register while its access register is set to zero, the results are undefined.

14. MICRO-PROCESSOR INTERFACE

A standard byte wide micro-processor interface (MPI) is used on all circuits with in the Spatial Decoder and Temporal Decoder... ..the parameter column. The actual specifications are shown in the respective columns min, max and units.

The DC operating conditions can be seen with reference to Table A.6.3. Here... ..the signal is present the maximum amount of time that this signal is available. The Units column gives the units of measurement used to describe the signals.

16. MPI WRITE TIMING

The general description of...Accordingly, the last picture will be held in the data buffer until a full swing buffer, but, by definition, the buffer will never fill. At some point, the machine will determine that an error condition exists... ..Consequently, the machine will not go into error recovery mode and will successfully continue to process the coded data.

A still further advantage of the use of a PICTURE(underscore)END token is that the serial pipeline processor will continue the processing of uninterrupted data. Through the use of a PICTURE(underscore)END token, the serial pipeline processor is configured to handle less than the expected amount of data and, therefore, continues processing. Typically, a prior art machine would stop itself because of an error condition. As previously... ..of the

Huffman decode and Video Demultiplexor know the number of blocks that it will process during each picture recovery cycle. When the correct number of blocks do not arrive from... Each of the state machines recognizes a FLUSH control token as information not to be processed. Accordingly, the FLUSH token is used to fill up all of the remaining empty parts of the coded data buffers and to allow a full set of information to be sent to the Huffman Decoder...less information than normally expected to decode the last picture. The Huffman decode circuit finishes processing the information contained in the last picture, and outputs this information through the DRAM interface... token, in accordance with the present invention, is used to pass through the entire pipeline processor and to ensure that the buffers are emptied and that other circuits are reconfigured to... END token, a padding word and a FLUSH token indicating to the serial pipeline processor that the picture processing for the current picture form is completed. Thereafter, the various state machines need reconfiguring to... FLUSH token resets each stage as it passes through, but-allows subsequent stages to continue processing. This prevents a loss of data. In other words, the FLUSH token is a variable... AFTER PICTURE

The STOP(underscore)AFTER(underscore)PICTURE function is employed to shut down the processing of the serial pipeline decompressing circuit at a logical point In Its operation. At this... a picture, the STOP(underscore)AFTER(underscore)PICTURE operation signals the end of all current processing.

22. MULTI(underscore)STANDARD - SEARCH MODE

Another feature of the present invention is the use... MODE control token which is used to reconfigure the input to the serial pipeline processor to look at the incoming bit stream. When the search mode is set, the Start... combination of control tokens, and DATA tokens along with the reconfiguration circuits, to provide similar processing.

The use of search mode in the present invention is convenient in many situations including... video disc. In general, a search mode is convenient when the user interrupts the normal processing of the serial pipeline at a point where the machine does not expect such an...be the case.

In brief, the Huffman Decoder 321 works in conjunction with the other units shown in Figure 27. These other units are the Parser State Machine 322, the inshifter 323, the Index to Data unit 324, the ALU 325, and the Token Formatter 326. As described previously, connection between these blocks is governed by a two wire interface. A more detailed description of how these units function is subsequently described herein in greater detail, the focus here is on particular aspects... control certain functions of the Index to Data 324 and ALU 325. Control of these units by the Huffman Decoder is necessary for proper decoding of block-level information. Having the... to be used by the Huffman Decoder for all three standards.

The Index to Data unit 324 performs the second part of the multi-part algorithm. This unit contains a look up table that provides the actual Huffman decoded data. Entries in the...by detecting these in the Huffman Decoder 321, rather than in the Index to Data unit 324.

This index number is then passed to the Index to Data unit 324. In essence, the Index to Data unit is a look-up table. In accordance with one aspect of the algorithm, the look... format that JPEG specifies for transferring an alternate JPEG table.

From the Index to Data unit 324, the decoded index number or other data is passed, together with the accompanying control... the entering data to ensure that the DATA tokens are of the correct size for processing. In fact, the token stream can be corrected in some situations if the error is... an order that is useful for the decompression circuits, but not for the particular display unit being used. When a block of data enters the Buffer Manager, the Buffer Manager supplies...the output of the Spatial Decoder or Temporal Decoder and re-format it for a computer or display system. The details of this formatting will vary between applications. In a simple...Token. The DATA Token can have as many bits as are necessary for carrying out processing at a particular place in the system. All other Tokens ignore the extra bits.

A.3.2 The DATA Token

The DATA Token carries data from one processing stage to the next. Consequently, the characteristics of this Token change as it passes through... will be sufficient to collect DATA Tokens and to detect a few Tokens that provide synchronization information (such as PICTURE(underscore)START). In this regard, see subsequent sections A.16, "Connecting... from the data stream. This provides an alternative to doing the configuration via the micro processor interface.

A.3.4 Description of Tokens

This section documents the Tokens which are implemented... 3.5.1. Note: JPEG requires a 2:1:1 structure for its macroblocks when processing 4:2:2 data. See Table A.3.5.

A.3.6 Special Token formats...either is low then the interface is taken to high impedance.

Note: on-chip data processing is not terminated when the DRAM interface is at high impedance. Therefore, errors will occur...decoded video's picture rate. Accordingly, this clock can be used to provide audio/video synchronization.

A.7.1 Spatial Decoder clock signals

The Spatial Decoder has two different (and potentially...in accordance with the present invention, must know what video standard is being input for processing. Thereafter, the system can accept either pre-existing Tokens or raw byte data which is...time a value is written into coded(underscore)data (7:0). Software is responsible for settling coded(underscore)extn to 0 before the last word of any Token is written to... 0). The start of this new DATA Token then passes into the Spatial Decoder for processing.

Each time a new 8 bit value is written to coded(underscore)data (7:0... Detector analyses data in the DATA Tokens bit serially. The Detector's normal rate of processing is one bit per clock cycle (of coded(underscore)clock). Accordingly, it will typically decode a byte of coded data every 8 cycles of coded(underscore)clock. However, extra processing cycles are occasionally required, e.g., when a non-DATA Token is supplied or when... Furthermore, this clock can be asynchronous to the main decoder(underscore)clock. Data transfer is synchronized to decoder(underscore)clock on-chip.

SECTION A.11 Start code detector

A.11.1...Code Detector. So, accessing these registers will be unreliable if the Start Code Detector is processing data. The user is responsible for ensuring that the Start Code Detector is halted before... Detector. In this case, the Tokens are passed through the Start Code Detector with no processing to other stages of the Spatial Decoder. These Tokens can only be inserted just before...result will be unpredictable if this is done when the Start Code Detector is actively processing data.

Discard all mode can be safely initiated after any of the Start Code Detector...start code non-alignment interrupt is suppressed.

In contrast, however, JPEG was designed for a computer environment where byte alignment is guaranteed. Therefore, marker codes should only be detected when byte... the other hand, was designed to meet the needs of both communications (bit serial) and computer (byte oriented) systems. Start codes in MPEG data should normally be byte aligned. However, the...result will be unpredictable if this is done when the Start Code Detector is actively processing data. So, before initiating a start code search, the Start Code Detector should be stopped so no data is being processed. The Start Code Detector is always in this condition if any of the Start Code...the spatial video decoding circuits (inverse modeler, quantizer and DCT). This second logical buffer allows processing time to include a spread so as to accommodate processing pictures having varying amounts of data.

Both buffers are physically held in a single off... if the buffers are full or empty.

As stated in A.13.1.1, the unit for all the above mentioned registers is a 512 bit block of data. Accordingly, the...until there is space in the buffer. If a buffer continues to be full, more processing stages "up steam" of the buffer will halt until the Spatial Decoder is unable to... ..converting coded data into Tokens started by the Start Code Detector. There are four main processing blocks in the Video Demux: Parser State Machine, Huffman decoder (including an ITOD), Macroblock counter... ..or state machine follows the syntax of the coded video data and instructs the other

units. The Huffman decoder converts variable length coded (VLC) data into integers. The Macroblock counter keeps track... ..In the present invention, picture dimensions are described to the Spatial Decoder in 2 different units: pixels and macroblocks. JPEG and MPEG both communicate picture dimensions in pixels. Communicating the dimensions... ..v and max(underscore)component(underscore)id specify the composition of the macroblocks (minimum coding units in JPEG). Each is a 2 bit register that can hold values in the range...supports some picture formats beyond those defined by JPEG and MPEG.

JPEG limits minimum coding units so that they contain no more than 10 blocks per scan. This limit does not apply to the Spatial Decoder since it can process any minimum coding unit that can be described by blocks(underscore)h(underscore)n, blocks(underscore)v(underscore)n... ..for 4:2:0 macroblocks (see Table A.14.8). However, the Spatial Decoder can process three other component macroblock structures, (e.g., 4:2:2.

A.14.5 Video events...of the Token buffer and the output of the Spatial Decoder.

There are three main units responsible for spatial decoding: the inverse modeler, the inverse quantizer and the inverse discrete cosine... ..At this point, the values in the DATA Tokens are quantized coefficients.

The inverse modelling process is the same regardless of the coding standard currently being used. No configuration is required...Inverse quantizer test registers

A.15.3 Inverse Discrete Cosine Transform

The inverse discrete transform processor of the present invention meets the requirements set out in CCITT recommendation H.261, the... ..with the requirements described in current draft revision of MPEG.

The inverse discrete cosine transform process is the same regardless of which coding standard is used. No, configuration by the user is required.

There are two events associated with the inverse discrete transform processor.

For a better understanding of the DCT and inverse DCT function the reader can examine...that a frame can contain.

Within an interleaved scan, data is organized into minimum coding units (MCUs) which are analogous to the macroblock used in MPEG and H.261. These MCUs... ..3.1 When to configure

The Temporal Decoder should only be configured when no data processing is taking place. This is the default state after reset is removed. The Temporal Decoder...

Claims: ...a video decoding system having an input, an output and a plurality of pipelined sequential processing stages between the input and the output, comprising :

a plurality of two-wire interfaces interconnecting... ..control and/or DATA tokens sequentially through said stages in the form of universal adaptation units for interfacing with said stages and interacting with selected stages, said two-wire interfaces each... ..a clock transition only when said sender is ready and said receiver is ready;

said processing stages comprising an image formatter receiving said tokens via a first said two-wire interface...
...allocating said buffers to said write address generator and said read address generator;

whereby said processing stages are afforded enhanced flexibility in configuration and processing.

2. The video decoding system according to claim 1, wherein said two-wire interfaces comprise...

6/K/14 (Item 14 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Country	Number	Kind	Date		
Type		Pub. Date		Kind	Text
Available Text		Language		Update	Word Count
Total Word Count (Document A)					
Total Word Count (Document B)					
Total Word Count (All Documents)					

Specification:

The present invention is directed to improvements in methods and apparatus for decompression which operates to decompress and/or decode a plurality of differently encoded input... ..of the well known standards known as JPEG, MPEG and H.261.

A serial pipeline processing system of the present invention comprises a single two-wire bus used for carrying unique... ..to a plurality of adaptive decompression circuits and the like positioned as a reconfigurable pipeline processor.

PRIOR ART

One prior art system is described in United States Patent No. 5,216,724. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and a host processor. The device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus.

United States Patent No, 4,785... ..a known quad tree data structure.

United States Patent No. 5,122,875 discloses an apparatus for encoding/decoding an HDTV signal. The apparatus includes a compression circuit responsive to high definition video source signals for providing hierarchically layered... ..compressed video data of relatively greater and lesser importance to image reproduction respectively. A transport processor, responsive to the high and low priority codeword sequences, forms high and low priority transport... ..United States Patent No. 5,168,356 discloses a video signal encoding system that includes apparatus for segmenting encoded video data into transport blocks for signal transmission. The ...in respective transport blocks.

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening. This is accomplished by an array transform processor such as that employed in a JPEG compression system. Blocks of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ..frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ..oscillators and the receiver can continuously receive each channel, then the receiver need not be synchronized with the transmitter. An FFT algorithm implements a fast discrete approximation to the continuous case in which the receiver synchronizes to the first frame and then acquires subsequent frames every frame period. The frame period... ..increasing the amount of data transmitted.

United States Patent No. 5,212,742 discloses an apparatus and method for processing video data for compression/decompression in real-time. The apparatus comprises a plurality of compute modules, in a preferred embodiment, for a total of four compute modules coupled in parallel. Each of the compute modules has a processor, dual port memory, scratch-pad memory, and an arbitration mechanism. A first bus couples the compute modules and host processor. Lastly, the device comprises a shared memory which is coupled to the host processor and to the compute modules with a second bus. The method handles assigning portions of the image for each of the processors to operate upon.

United States Patent No. 5,231,484 discloses a system and method... ..MPEG standards. Included are three cooperating components or subsystems that operate to variously adaptively pre-process the incoming digital motion video sequences, allocate bits to the pictures in a sequence, and... ..States Patent No. 5,267,334 discloses a method of removing frame redundancy in a computer system for a sequence of moving images. The method comprises detecting a first scene change... ..facing" keyframe or intraframe, and it is normally present in CCITT compressed video data. The process then comprises generating at least one intermediate compressed frame, the at least one intermediate compressed frame containing difference information from the first image for at least one image following...change, known as a "backward-facing" keyframe. The first keyframe and the at least one intermediate compressed frame are linked for forward play, and the second keyframe and the intermediate compressed frames are linked in reverse for reverse play. The intraframe may also be used... ..of complete scene information.

United States Patent No. 5,276,513 discloses a first circuit apparatus, comprising a given number of prior-art image-pyramid stages, together with a second circuit apparatus, comprising the same given number of novel motion-vector stages, perform cost-effective hierarchical motion analysis (HMA) in real-time, with minimum system processing delay and/or employing minimum system processing delay and/or employing minimum hardware structure. Specifically, the first and second circuit apparatus, in response to relatively high-resolution image data from an ongoing input series of successive... ..a relatively high frame rate (e.g., 30 frames per second), derives, after a certain processing-system delay, an ongoing output series of successive given pixel-density vector-data frames that... ..of successive image frames.

United States Patent No. 5,283,646 discloses a method and apparatus for enabling a real-time video encoding system to accurately deliver the desired number of... ..desired bit allocations.

The article, Chong, Yong M., A Data-Flow Architecture for Digital Image Processing, Wescon Technical Papers: No. 2 Oct./Nov. 1984, discloses a real-time signal processing system specifically designed for image processing. More particularly, a token based data-flow architecture is disclosed wherein the tokens are of... ..width having a fixed width address field. The system contains a plurality of identical flow processors connected in a ring fashion. The tokens contain a data field, a control field and a tag. The tag field of the token is further broken down into a processor address field and an identifier field. The processor address field is used to direct the tokens to the correct data-flow processor, and the identifier field is used to label the data such that the data-flow processor knows what to do with the data. In this way, the identifier field acts as an instruction for the data-flow processor. The system directs each token to a specific data-flow processor using a module number (MN). If the MN matches the MN of the

particular stage... ..to locate the decoder in the preceding stage in order to pre-decode complex decoding processing and to alleviate critical path problems in the logic circuit. The elastic nature of the...of block signal in most cases.

United States Patent No. 4,903,018 discloses a process and data processing system for compressing and expanding structurally associated multiple data sequences. The process is particular to data sets in which an analysis is made of the structure in... ..data series on the basis of the order number of these data elements. The data processing system for performing the processes includes a storage matrix (26) and an index storage (28) having line addresses of the... ..the final actual video.

United States Patent No. 5,060,242 discloses an image signal processing system DPCM encodes the signal, then Huffman and run length encodes the signal to produce... ..tightly packed without gap for efficient transmission without loss of any data. The tightly packed apparatus has a barrel shifter with its shift modulus controlled by an accumulator receiving code word... ..OR gate is connected to the shifter, while a register is connected to the gate. Apparatus for processing a tightly packed and decorrelated digital signal has a barrel shifter and accumulator for unpacking... ..an inverse DCPM decoder,

United States Patent No. 5,168,375 discloses a method for processing a field of image data samples to provide for one or more of the functions of decimation, interpolation, and sharpening is accomplished by use of an array transform processor such as that employed in a JPEG compression system. Block of data samples are transformed by the discrete even cosine transform (DECT) in both the decimation and interpolation processes, after which the number of frequency terms is altered. In the case of decimation, the... ..frequency domain, there is provided an inverse transformation resulting in a set of blocks of processed data samples. The blocks are overlapped followed by a savings of designated samples, and a... ..kernel matrix.

United States Patent No. 5,231,486 discloses a high definition video system processes a bitstream including high and low priority variable length coded Data words. The coded Data... ..packed High Priority Data and packed Low Priority Data by means of respective data packing units. The coded Data is continuously applied to both packing units. High Priority and Low Priority Length words indicating the bit lengths of high priority and... ..States Patent No. 5,287,178 discloses a video signal encoding system includes a signal processor for segmenting encoded video data into transport blocks having a header section and a packed data section. The system also includes reset control apparatus for releasing resets of system components, after a global system reset, in a prescribed non-simultaneous phased sequence to enable signal processing to commence in the prescribed sequence. The phased reset release sequence begins when valid data...United States Patent No. 5,142,380 to Sakagami et al. discloses an image compression apparatus suitable for use with still images such as those formed by electronic still cameras using... ..and Q.

United States Patent No. 5,193,002 to Guichard et al. disclosed an apparatus for coding/decoding image signals in real time in conjunction with the CCITT standard H.261. A digital signal processor carries out direct quantization and reverse quantization.

United States Patent No. 5,241,383 to Chen et al. describes an apparatus with a pseudo-constant bit rate video coding achieved by an adjustable quantization parameter. The... ..relates to an improved pipeline system having an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along the pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the processing stages in the pipeline and interacting with selected stages in the pipeline for control data and/or combined control-data functions among the processing stages, so that the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing. In accordance with the invention, the processing stages may be configurable in response to recognition of at least one token. One of the processing stages may be a Start Code Detector which receives the input and generates and/or... ..and resetting the system, and a CODING(underscore)STANDARD token

for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards. The present invention... ..data and having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video pictures of varying data size can be controlled. Also in accordance with the invention, a processing stage receives the input data stream, the stage including means for recognizing specified bit stream patterns, whereby the processing stage facilitates random access and error recovery, The invention may also include a means for... ..invention also includes an inverse modeller stage, an inverse discrete cosine transform stage, and a processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for processing data.

In addition, the present invention relates to an improved pipeline system having a Huffman...pipeline stage that incorporates a two-wire transfer control and also shows two consecutive pipeline processing stages with the two-wire transfer control;

Figures. 5a and 5b taken together depict one... ..shown in Figures. 8a and 8b.

Figure 10 is a block diagram of a reconfigurable processing stage;

Figure 11 is a block diagram of a spatial decoder;

Figure 12 is a... ..decoder including the prediction filters;

Figure 18 is a pictorial representation of the prediction filtering process;

Figure 19 shows a generalized representation of the macroblock structure;

Figure 20 shows a generalized... ..buffer;

Figure 25 is a pictorial diagram illustrating prediction data offset from the block being processed;

Figure 26 is a pictorial diagram illustrating prediction data offset by (1,1);

Figure 27...in general terms, the present invention provides an input, an output and a plurality of processing stages between the input and the output, the plurality of processing stages being interconnected by a two-wire interface for conveyance of tokens along a pipeline, and control and/or DATA tokens in the form of universal adaptation units for interfacing with all of the stages in the pipeline and interacting with selected stages in the pipeline for control, data and/or combined control-data functions among the processing stages, whereby the processing stages in the pipeline are afforded enhanced flexibility in configuration and processing.

Each of the processing stages in the pipeline may include both primary and secondary storage, and the stages in... ..The tokens in the pipeline are dynamically adaptive and may be position dependent upon the processing stages for performance of functions or position independent of the processing stages for performance of functions.

In a pipeline machine, in accordance with the invention, the... ..altered by interfacing with the stages, and the tokens may interact with all of the processing stages in the pipeline or only with some but less than all of said processing stages. The tokens in the pipeline may interact with adjacent processing stages or with non-adjacent processing stages, and the tokens may reconfigure the processing stages. Such tokens may be position dependent for some functions and position independent for other... ..be Huffman coded.

In the improved pipeline machine, the tokens may be generated by a processing stage. Such pipeline tokens may include data for transfer to the processing stages or the tokens may be devoid of data. Some of the tokens may be identified as DATA tokens and provide data to the processing stages in the pipeline, while other tokens are identified as control tokens and only condition the processing stages in the pipeline, such conditioning including reconfiguring of the processing stages. Still other tokens may provide both data and conditioning to the processing stages in the pipeline. Some of said tokens may identify coding standards to the processing stages in the pipeline,

whereas other tokens may operate independent of any coding standard among the processing stages. The tokens may be capable of successive alteration by the processing stages in the pipeline.

In accordance with the invention, the interactive flexibility of the tokens in cooperation with the processing stages facilitates greater functional diversity of the processing stages for resident structure in the pipeline, and the flexibility of the tokens facilitates system... or alteration. The tokens may be capable of facilitating a plurality of functions within any processing stage in the pipeline. Such pipeline tokens may be either hardware based or software based... system bandwidth in the pipeline. The tokens may provide data and control simultaneously to the processing stages in the pipeline.

The invention may include a pipeline processing machine for handling plurality of separately encoded bit streams arranged as a single serial bit... and for passing unrecognized control tokens along the pipeline, and a reconfigurable decode and parser processing means responsive to a recognized control token for reconfiguring a particular stage to handle an... be a pipeline system and the Start Code Detector may be positioned as the first processing stage in the pipeline.

The present invention also provides, in a system having a plurality of processing stages, a universal adaptation unit in the form of an interactive interfacing token for control and/or data functions among the processing stages, the token being a PICTURE(underscore)START code token for indicating that the start... The token may also be a CODING(underscore)STANDARD token for conditioning the system for processing in a selected one of a plurality of picture compression/decompression standards.

The CODING(underscore... standard as JPEG, and/or any other appropriate picture standard. At least some of the processing stages reconfigure in response to the CODING(underscore)STANDARD token.

One of the processing stages in the system may be a Huffman decoder and parser and, upon receipt of... Data stage, and the parser stage may send an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... video data, having a Huffman decoder, an index to data (ITOD) stage, an arithmetic logic unit (ALU), and a data buffering means immediately following the system, whereby time spread for video... be controlled.

The system may include spatial decoder having a two-wire interface interconnecting processing stages, the interface enabling serial processing for data and parallel processing for control.

As previously indicated, the system may further include a ROM having separate stored... of a plurality of picture standards, the programs being selectable by a token to facilitate processing for a plurality of different picture standards.

The spatial decoder system also includes a token... decoding stage and a parser stage for sending an instruction to the Index to Data Unit to select tables needed for a particular identified coding standard, the parser stage indicating whether... The present invention also provides a pipeline system having an input data stream, and a processing stage for receiving the input data stream, the stage including means for recognizing specified bit... whereby said stage facilitates random access and error recovery. In accordance with the invention, the processing stage may be a start code detector and the bit stream patterns may include start... token and padding insures uniformity of word size. In accordance with the invention, a reconfigurable processing stage may be provided as a spatial decoder and the padding means adds to picture... that if the DATA token has less than the predetermined length, the padder circuit adds units of data to the DATA token until the predetermined length is achieved. A bypass circuit... tokens into a buffer, having a second predetermined width.

The invention also provides an apparatus for providing a time delay to a group of compressed pictures, the pictures corresponding to... and capable of delaying the words of data, is in communication with a control circuit intermediate the counter circuit and the inverse modeller circuit, the control circuit also communicating with the... inverse modeller stage and an inverse discrete cosine transform stage, the improvement characterized by a

processing stage, positioned between the inverse modeller stage and the inverse discrete cosine transform stage, responsive to a token table for

processing data.

In accordance with the invention, the token may be a QUANT(underscore)TABLE token for causing the processing stage to generate a quantization table.

The present invention also provides a Huffman decoder for... ..of bits used to represent an item of data.

DECODER: An embodiment of a decoding process.

DECODING (PROCESS): The process defined in this specification that reads an input coded bitstream and produces decoded pictures or... ..the same order in which they were presented at the input of the encoder.

ENCODING (PROCESS): A process, not specified in this specification, that reads a stream of input pictures or audio samples...to provide an estimate of the pel value or data element currently being decoded.

RECONFIGURABLE PROCESS STAGE (RPS): A stage, which in response to a recognized token, reconfigures itself to perform various operations.

SLICE: A series of macroblocks.

TOKEN: A universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions... ..indicates that the corresponding stage holds valid data, i.e., data that is to be processed in one of the pipeline stages. After processing (which may involve nothing more than a simple transfer without manipulation of the data) valid... ..present invention may be used with any number of pipeline stages. Furthermore, data may be processed in more than one stage and the processing time for different stages can differ.

In addition to clock and data signals (described below... ..other system. For example, the last pipeline stage may pass its data on to subsequent processing circuitry. The ACCEPT signal, which is illustrated as the lower of the two lines connecting... ..the minimum disturbance possible to other pipeline stages. Succeeding pipeline stages are allowed to continue processing and, therefore, this means that gaps open up in the stream of data following the... ..The data in the pipeline is encoded such that many different types of data are processed in the pipeline. This encoding accommodates data packets of variable size and the size of...the other hand, it may generate itself, all or part of the data to be processed in the pipeline. Indeed, as is explained below, a "stage" may contain arbitrary processing circuitry, including none at all (for simple passing of data) or entire systems (for example... ..values zero and 255 may not be used.

If such a picture were to be processed in a pipeline built in the practice of the present invention, then one of these... ..data must not be written over since it is data that must be saved for processing or use in a downstream device e.g., a pipeline stage, a device or a... ..connected to the pipeline upstream contains data D4 that is to be transferred into and processed in the pipeline. Stages B, D and E, in addition ...pipeline, in accordance with the preferred embodiments of the present invention, to "fill up" empty processing stages is highly advantageous since the processing stages in the pipeline thereby become decouple from one another. In other words, even though... ..data can be transferred into the pipeline and between stages even when one or more processing stages is blocked.

In the embodiment shown in Fig. 1, it is assumed that the...propagate all the way back to the beginning of the pipeline if there is some intermediate stage that is able to accept new data.

In the embodiment illustrated In Fig. 1...has been mentioned. It is to be further understood that each pipeline stage may also process the data it has received arbitrarily before passing it between its internal storage elements or... ..the

portion of the pipeline that contains input and output storage elements and that arbitrarily processes data stored in its storage elements.

Furthermore, the "device" downstream from the pipeline Stage F...valid data, but also when a stage requires more than one clock phase to finish processing its data. This also can occur when it creates valid data in one or both...
...control the passage of data between adjacent storage elements. The VALID signal may also be processed in an analogous manner.

A great advantage of the two-wire interface (one wire for... ..In addition, two extra latches and a small number of gates are preferably added to process the ACCEPT and VALID signals that are associated with the data latches in each half... ..application so requires. The interface in accordance with this embodiment can also be used to process analog signals.

As discussed previously, while other conventional timing arrangements may be used, the interface... ..circuit B1, which may be provided to convert output data from input latch LDIN into intermediate data, which is then later loaded in an output data latch LDOUT, which comprises the... ..is connected either directly as an input to the validation output latch LVOUT, or via intermediate logic devices or circuits that may alter the signal.

Similarly, the output validation signal QVOUT... ..to the input of the validation input latch QVIN of the following stage, or via intermediate devices or logic circuits, which may alter the validation signal. This output QVIN is also...word.

Preferred Data Structure - "tokens"

In the sample application shown in Fig. 4, each stage processes all input data, since there is no control circuitry that excludes any stage from allowing... ..are connected together in a relatively simple configuration. The simplest configuration is a pipeline of processing steps. For example, in the one shown in Fig. 1. The use of tokens, however left to right in the diagram. Data enters the machine and passes into processing Stage A. This may or may not modify the data and it then passes the... ..advantage of the tokens is their ability to achieve this kind of communication. Since any processing stage that does not recognize a token simply passes it on unaltered to the next... ..is transmitted along with the address and data fields in each token so that a processing stage can pass on a token (which can be of arbitrary length) without having to... ..be the first word of a new token.

Note that although the simple pipeline of processing stages is particularly useful, it will be appreciated that tokens may be applied to more complicated configurations of processing elements. An example of a more complicated processing element is described below.

It is not necessary, in accordance with the present invention, to... ..has extension bits. An example of this is a token that activates a stage that processes video quantization values stored in a quantization table (typically a memory device). For example, a...turn, is of great importance in video data pipeline systems since it ensures that all processing stages can be continuously running at full bandwidth.

In accordance with the present invention, in... ..some other chips in the set. This is advantageous both from the perspective of a customer and from that of a chip manufacturer. Even if modifications mean that all chips are...the end of a token (and hence the start of the next token) to be processed correctly (Including simple non-manipulative transfer), even if the token is not recognized by the... ..block diagram of a pipeline stage whose function is as follows. If the stage is processing a predetermined token (known in this example as the DATA token), then it will duplicate... ..the address field of the DATA token. If, on the other hand, the stage is processing any other kind of token, it will delete every word. The overall effect is that... ..respective output signals:

In the duplication stage, the output from the data latch LDIN forms intermediate data referred to as MID(underscore)DATA. This intermediate data word is loaded into the data output latch LDOUT only when an intermediate acceptance signal (labeled "MID(underscore)ACCEPT" in Fig. 8a) is set HIGH.

The portion of... data. These include a "DATA(underscore)TOKEN" signal that indicates that the circuitry is currently processing a valid DATA Token, and a NOT(underscore)DUPLICATE signal which is used to control duplication of data. When the circuitry is processing a DATA Token, the NOT(underscore)DUPLICATE signal toggles between a HIGH and a LOW... the token to be duplicated once (but no more times). When the circuitry is not processing a valid DATA Token then the NOT(underscore)DUPLICATE signal is held in a HIGH state. Accordingly, this means that the token words that are being processed are not duplicated.

As Fig. 8a illustrates, the upper six bits of 8-bit intermediate data word and the output signal QI1 from the latch LI1 form inputs to a... explained further below.

Latch LO1 performs the function of latching the last value of the intermediate extension bit (labeled "MID(underscore)EXTN" and as signal S4), and it loads this value... and the DATA(underscore)TOKEN signal will become "0", indicating that the circuitry is not processing a DATA token.

If QI1 is "0" and SO is "0", thereby indicating a DATA... phase and the DATA(underscore)TOKEN signal will become "1", indicating that the circuitry is processing a DATA token.

The NOT(underscore)DUPLICATE signal (the output signal QO3) is similarly loaded... LVOUT at the same time that MID(underscore)DATA is loaded into LDOUT and the intermediate extension bit (signal S4) is loaded into LEOUT. Signal S5 is also combined with the... above. This has the effect that all tokens except the one that causes the duplication process will be deleted from the token stream, since a device connected to the output terminals (OUTDATA, OUTEXTN and OUTVALID) will not recognize these token words as valid data.

As before... and is duplicated.

Referring now more particularly to Figure 10, there is shown a reconfigurable process stage in accordance with one aspect of the present invention.

Input latches 34 receive an... the input latches 34 is passed as a first input over line 35 to a processing unit

36. A first output from the token decode subsystem 33 is passed over line 37 as a second input to the

processing unit 36. A second output from the token decode 33 is passed over line 40 to an action identification unit 39. The action identification unit 39 also receives input from registers 43 and 44 over line 46. The registers 43... is determined by the history of tokens previously received. The output from the action identification unit 38 is passed over line 38 as a third input to the processing unit 36. The output from the processing unit 36 is passed to output latches 41. The output from the output latches 41 is... decoder 56 is passed over line 63 as an input to an Index to Data Unit (ITOD) 64. The Huffman decoder 56 and the ITOD 64 work together as a single logical unit. The output from the ITOD 64 is passed over line 65 to an arithmetic logic unit (ALU) 66. A first output from the ALU 66 is passed over line 67 to... blocks 133.

Referring to Figure 14b, in the JPEG and H.261 standards, the Common Intermediate Format (CIF) is used, wherein a picture 141 is encoded as 6 rows each containing... in a zigzag direction indicated by the arrow 144. The GOBs 142 are, in turn, processed row-by-row, left-to-right in each row.

Referring now to Figure 14c, it... in accordance with the practice of the present invention. A first picture 161 to be processed contains a first PICTURE(underscore)START token 162, first-picture information of indeterminate length 163, and a first PICTURE(underscore)END token 164. A second picture 165 to be processed contains a second PICTURE(underscore)START token 166, second picture information of indeterminate length 167... tokens 162 and 166 indicate the start of the pictures 161 and 165 to the processor. Likewise, the PICTURE(underscore)END tokens 164 and 168 signify the end of the pictures 161 and 165 to the processor. This allows the processor to process picture information 163 and 167 of variable lengths.

Referring to Figure 17, a split 171...Video Formatter (not shown in Figure 17).

Referring now to Figure 18, the prediction filtering process is illustrated. A forward picture 201 is passed over line 202 as a first input... ..the right of the value decode shift register 230, as indicated by area 231. This process eliminates overlapping start code images, as discussed below. A first output from the value decode... ..Code Detector. The Start Code Detector then receives a first data value image 244. Before processing the first data value image 244, the Start Code Detector may detect a second start... ..image 244 at a length 246. If this occurs, the Start Code Detector does not process the first data value image 244, and instead receives and processes a second data value image 247.

Referring now to Figure 22, a flag generator 251...line 1 of Table 600, whenever a "sequence start" image is received during H.261 processing or a "picture start" image is received during MPEG processing, the entire group of four control tokens is generated, each followed by its corresponding data...Picture Decoding

3. Motion Picture Decompression

4. RAM Memory Map

5. Bitstream Characteristics

6. Reconfigurable Processing Stage

7. Multi-Standard Coding

8. Multi-Standard Processing Circuit-2nd Mode of Operation

9. Start Code Detector

10. Tokens

11. DRAM Interface

12... ..described herein in greater detail) and reformatting this output for use, including display in a computer or other display systems, including a video display system. Implementation of this formatting varies significantly... ..the Spatial Decoder circuits.

The Spatial Decoder of the present invention performs all the required processing within a single picture. This reduces the redundancy within one picture.

The Temporal Decoder reduces... ..modeller 75, the inverse zig-zag 81 and the inverse DCT 83. The standard independent units within the Huffman decoder and parser include the ALU 66 and the token formatter 71.

Referring now to Figure 12, the standard-independent units include the DRAM interface 100, the fork 91, the FIFO register 96, the summer 98 and the output selector 106. The standard dependent units are the address generator 94, which is different in H.261 and in MPEG, and...much of the operation is very similar between the three different compression standards.

The next unit is the state machine 68 (Figure 11) located within the Huffman decoder and parser. Here... ..The same holds true for JPEG, which is a third, completely independent program.

The next unit is the Huffman decoder 56 which functions with the index to data unit 64. Those two units cooperate together to perform the Huffman decoding. Here, the algorithm that is used for Huffman... ..to the Huffman decoder at different times consistent with the standard in operation.

The last unit on the chip that is dependent on the compression standard is the inverse quantizer 79...an H.261 group of blocks and an MPEG slice. When H.261 data is processed after the Start Code Detector, each group of blocks is preceded by a slice(underscore... ..these standards have totally different sets of tables.

As previously indicated, most of the system units are compression standard independent. If a unit is standard independent, and such units need not remember what CODING(underscore)STANDARD is being processed. All of the units that are standard dependent remember the compression standard as the CODING(underscore)STANDARD token flows... ..CODING(underscore)STANDARD tokens at the Start Code Detector that is positioned as the first unit in the pipeline, this change of compression standard is readily handled. The token says a... ..found in the standard, i.e. from the bitstream into a prediction mode token. This processing is performed by the Huffman decoder and parser state machine, where it is easy to... ..to that token. By having these tokens and using them appropriately, the design of other units in the machine is simplified. Although there may be some complications in the program, benefits... ..a first encoded signal (the MPEG or H.261 encoded video signal) in a pipeline processing system. The Temporal Decoder is not needed for JPEG decoding.

In this regard, the invention... ..the use of a single pipeline decoder and decompression system. The decoding-and decompression pipeline processor is organized on a unique and special configuration which allows the handling of the multistandard... ..video signals through the use of techniques all compatible with the single pipeline decoder and processing system. The Spatial Decoder is combined with the Temporal Decoder, and the Video Formatter is... ..with only still pictures. The compression standard independent Spatial Decoder performs all of the data processing within the boundaries of ...to the multi-standard, configurable Video Formatter, which then provides an output to the display terminal. In a first sequence of similar pictures, each decompressed picture at the output of the... ..of control tokens and DATA tokens, in combination with a plurality of sequentially-positioned reconfigurable processing stages selected and organized to act as a standard-independent, reconfigurable-pipeline-processor.

With regard to JPEG decoding, a single Spatial Decoder with no off chip DRAM can... ..video. Accordingly, signals carried by DATA tokens pass directly through the Temporal Decoder without further processing when the Temporal Decoder is configured for a JPEG operation.

Another aspect of the present... ..for subsequent use in temporal decoding of subsequent pictures.

Generally, the Temporal Decoder performs the processing between pictures either-earlier and/or later in time with reference to the picture currently... ..is distributed among several areas of DRAM in the sense that the decompressed output information, processed by the Spatial Decoder, is stored in other DRAM registers by other random access memories...first decoder circuit (the Spatial Decoder) directly to the Video Formatter for handling without signal processing delay.

The Temporal Decoder also reorders the blocks of picture data for display by a... ..from a selection of pictures which have arrived earlier or later than the picture under processing. When a picture is described in this context, it may mean any one of the... ..2. The result, i.e., the final decoded picture resulting from the addition of a process step performed by the decoder;

3. Previously decoded pictures read from the DRAM; and

4... ..START token and a subsequent PICTURE(underscore)END token.

After the picture data information is processed by the Temporal Decoder, it is either displayed or written back into a picture memory location. This information is then kept for further reference to be used in processing another different coded data picture.

Re-ordering of the MPEG encoded pictures for visual display...used to encode a referenced picture of a picture might be identified as being one unit long, another picture might be a number of units long, while still a third picture could be a fraction of that unit.

None of the existing standards (MPEG 1.2, JPEG, H.261) define a way of... ..picture rate, whereas the Video Formatter can handle a variable input picture rate.

6. RECONFIGURABLE PROCESSING STAGE

Referring again to Figure 10, the reconfigurable processing stage (RPS) comprises a token decode circuit 33 which is employed to receive the tokens... input latches 34. The output of the token decode circuit 33 is applied to a processing unit 36 over the two-wire interface 37 and an action identification circuit 39. The processing unit 36 is suitable for processing data under the control of the action identification circuit 39. After the processing is completed, the processing unit 36 connects such completed signals to the output, two-wire interface bus 40 through output... token decode circuit 33 are applied simultaneously to the action identification circuit 39 and the processing unit 36. The action identification function as well as the RPS is described in further detail... not standard independent circuits. The data flows through the token decode circuit 33, through the

processing unit 36 and onto the two-wire interface circuit 42 through the output latches 41. If... wire interface 42 through the output circuit 41. The present invention operates as a pipeline processor having a two-wire interface for controlling the movement of control tokens through the pipeline... time, the token decode circuit 33 provides a proper flag or index signal to the processing unit 36 to alert it to the presence of the token being handled by the action identification circuit 39. Control tokens may also be processed.

A more detailed description of the various types of tokens usable in the present invention... standard now passing through the state machine shown with reference to Figure 10.

Similarly, the processing unit 38 which is under the control of the action identification circuit 39 is now ready to process the information contained in the data fields of the DATA token when it is appropriate... action identification circuit 39 and is immediately followed by a DATA token which is then processed by the processing unit 36. The control token exits the output latches circuit 41 over the output two-wire interface 42 immediately preceding the DATA token which has been processed within the processing unit 36.

In the present invention, the action identification circuit, 39, is a state machine holding... show that the action can also be affected by the token that is currently being processed by the token decode circuit 33.

In general, there is shown token decoding and data processing in accordance with the present invention. The data processing is performed as configured by the action identification circuit 39. The action is affected by... information stored from previously decoded tokens in registers 43 and 44, the current token under processing, and the state and history information that the action identification unit 39 has itself acquired. A distinction is thereby shown between Control tokens and DATA tokens.

In any RPS, some tokens are viewed by that RPS unit as being Control tokens in that they affect the operation of the RPS presumably at... are viewed by the RPS as DATA tokens. Such DATA tokens contain information which is processed by the RPS in a way that is determined by the design of the particular... view of the same token. Some of the tokens might be viewed by one RPS unit as DATA Tokens while another RPS unit might decide that it is actually a Control Token. For example, the quantization table information... into a token called a quantization table token (QUANT(TABLE)) which goes down the processing pipeline. As far as that machine is concerned, all of that was data; it was... sort of data into another sort of data, which is clearly a function of the processing performed by that portion of the machine. However, when that information gets to the inverse... present. This information is viewed as control information, and then that control information affects the processing that is done on subsequent DATA tokens because it affects the number that you multiply... important feature of the invention is that each of the stages of circuitry has the processing capability within it to be able to perform the necessary operations for each of the... operations are to be performed at a given time, come as tokens. There is one processing element that differs between the different stages to provide this capability. In the state machine... standard is and it looks up the parameters that it needs to apply to the processing elements in order to perform a proper operation. For example, the inverse quantizer will look... is set to 1 for a particular compression standard, and will apply that to its processing circuitry.

In a similar sense the Huffman decoder 56 has a number of tables within... MPEG video standard or the JPEG video standard. These three compression coding standards specify similar processes to be done on the arriving data, but the structure of the datastreams is different... token stream embodying the current coding standard. The control tokens are passed through the pipeline processor, and are used, i.e., decoded, in the state machines to which they are relevant... this regard, the DATA Tokens are treated in the same fashion, insofar as they are processed only in the state machines that are configurable by the control tokens into processing such DATA Tokens. In the remaining state machines, they pass through unchanged.

More specifically, a... signals. The remaining portions of the token are used to indicate and identify the internal processing control function which is standard for all of the datastreams passing through the pipeline processor. In one form of the invention, the token extension is used to carry the current... accompanying data. As previously discussed, this information is utilized in the system to reconfigure the processing stage used to perform the function required by the various standards created for that purpose... picture number as indicated by the value.

The system also includes a multi-stage parallel processing pipeline operating under the principles of the two-wire interface previously described. Each of the... the token presently entering the state machine into the action identification circuit 39 or the processing unit 36, as appropriate. The processing unit has been previously reconfigured by the next previous control token into the form needed for handling the current coding standard, which is now entering the processing stage and carried by the next DATA token. Further, in accordance with this aspect of the invention, the succeeding state machines in the processing pipeline can be functioning under one coding standard, i.e., H.261, while a previous... tokens required to decode a number of coding standards with a fixed number of reconfigurable processing stages. More specifically, the PICTURE(underscore)END control token is employed because it is important... standard machine, it is necessary to create additional control tokens within the multi-standard pipeline processing machine which will then indicate which one of the standard decoding techniques to use. Such... and to push the current picture through the decoder to the display.

8. MULTI-STANDARD PROCESSING CIRCUIT - SECOND MODE OF OPERATION

A compression standard-dependent circuit, in the form of the... of the Start Code Detector will subsequently be discussed in further detail, as will the process of starting up of the decoder.

The aforementioned description has been concerned primarily with the... the data which immediately follows according to the standard. However, in the multi-standard pipeline processing system of the present invention, where compatibility is required for multiple standards, the system has... signals, including flag signals, are generated by each state machine to handle some of the processing within that state machine. Values carried in the standards can be used to access machine... its contents must be removed from the two wire interface to ensure that no further processing takes place using these 3 bytes. The decode register is emptied, and the value decode... 10. TOKENS

In the practice of the present invention, a token is a universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions and is adapted for use with a reconfigurable processing stage (RPS) which is a stage, which in response to a recognized token, reconfigures itself to perform various operations.

Tokens may be either position dependent or position independent upon the processing stages for performance of various functions. Tokens may also be metamorphic in that they can be altered by a processing stage and then passed down the pipeline for performance of further functions. Tokens may interact... other functions, and the specific interaction with a stage may be conditioned by the previous processing history of a stage.

A PICTURE(underscore)END token is a way of signaling the... through a fixed size, fixed width buffer.

The present invention is directed to a pipeline processing system which has a variable configuration which uses tokens and a two-wire system. The... do not use control tokens.

The control tokens are generated by circuitry within the decoder processor and emulate the operation of a number of different type standard-dependent signals passing into the serial pipeline processor for handling. The technique used is to study all the parameters of the multi-standards that are selected for processing by the serial processor and noting 1) their similarities, 2) their dissimilarities, 3) their needs and requirements and 4) selecting the correct token function to effectively process all of the standard signals sent into the serial processor. The functions of the tokens are to emulate the standards. A control token function is... ..the standard dependent signals and as an element to transmit control information through the pipeline processor

In prior art system, a dedicated machine is designed according to well-known techniques to... ..tokens provide and make a sensible format for communicating information through the decompression circuit pipeline processor. In the design selected hereinafter and used in the preferred embodiment, each word of a... ..However, this is not a limitation on the invention, but on the magnitude of the processing steps elected to be accomplished by use of these tokens. It is to be noted... ..bit address for use in accessing the random access memories used throughout this serial decompression processor. This provides an additional degree of variability that facilitates a broad range of versatility.

As previously described, the DATA token carries data from one processing stage to the next. Consequently, the characteristics of this token change as it passes through...longest number of data bits because it needs to provide the most information to the

processing unit so that it can start the decompression with as much information as possible. Words which...to receive an address, it waits for the address generator to supply a valid address, processes that address and then sets the accept line high for one clock period. Thus, it... ..be read. This signal passes between two asynchronous clock regimes and, therefore, passes through three synchronizing flip flops.

Provided RAM2 312 is empty, the next item of data to arrive on...interesting.

In general, prediction data will be offset from the position of the block being processed as specified in the motion vectors in x and y. Thus, the block of data... ..address, 9. Data is read from this address and the x value is incremented. The process is repeated until the x value reaches its stop value, at which point, the y... ..is read, the x value is again incremented until it reaches its stop value. The process is repeated until both x and y values have reached their stop values. Thus, the...invention, is that additional information must be provided to the prediction filters to indicate what processing is required on the data. This consists of the following:

a "last byte" signal indicating... ..bit 0) is incremented and the x address (3 LSBS) is reset to zero. This process is repeated until 64 bytes have been read. With a 16 or 32 bit wide...register while its access register is set to zero, the results are undefined.

14. MICRO-PROCESSOR INTERFACE

A standard byte wide micro-processor interface (MPI) is used on all circuits with in the Spatial Decoder and Temporal Decoder... ..the parameter column. The actual specifications are shown in the respective columns min, max and units.

The DC operating conditions can be seen with reference to Table A.6.3. Here... ..the signal is present the maximum amount of time that this signal is available. The Units column gives the units of measurement used to describe the signals.

16. MPI WRITE TIMING

The general description of...a PICTURE(underscore)END token is decoded and forces the data in the coded data buffers to be applied to the Huffman decoder and video demultiplexor, the final picture can be... Consequently, the machine will not go into error recovery mode and will successfully continue to process the coded data.

A still further advantage of the use of a PICTURE(underscore)END token is that the serial pipeline processor will continue the processing of uninterrupted data. Through the use of a PICTURE(underscore)END token, the serial pipeline processor is configured to handle less than the expected amount of data and, therefore, continues processing. Typically, a prior art machine would stop itself because of an error condition. As previously... of the Huffman decode and Video Demultiplexor know the number of blocks that it will process during each picture recovery cycle. When the correct number of blocks do not arrive from... Each of the state machines recognizes a FLUSH control token as information not to be processed. Accordingly, the FLUSH token is used to fill up all of the remaining empty parts... Huffman Decoder and Video Demultiplexor. In this way, the FLUSH token is like padding for buffers.

The Token Decoder in the Huffman circuit recognizes the FLUSH token and ignores the pseudo...less information than normally expected to decode the last picture. The Huffman decode circuit finishes processing the information contained in the last picture, and outputs this information through the DRAM interface... token, in accordance with the present invention, is used to pass through the entire pipeline processor and to ensure that the buffers are emptied and that other circuits are reconfigured to... (underscore)END token, a padding word and a FLUSH token indicating to the serial pipeline processor that the picture processing for the current picture form is completed. Thereafter, the various state machines need reconfiguring to... FLUSH token resets each stage as it passes through, but-allows subsequent stages to continue processing. This prevents a loss of data. In other words, the FLUSH token is a variable... AFTER PICTURE

The STOP(underscore)AFTER(underscore)PICTURE function is employed to shut down the processing of the serial pipeline decompressing circuit at a logical point in its operation. At this... a picture, the STOP(underscore)AFTER(underscore)PICTURE operation signals the end of all current processing.

22. MULTI(underscore)STANDARD - SEARCH MODE

Another feature of the present invention is the use... (underscore)MODE control token which is used to reconfigure the input to the serial pipeline processor to look at the incoming bit stream. When the search mode is set, the Start... combination of control tokens, and DATA tokens along with the reconfiguration circuits, to provide similar processing.

The use of search mode in the present invention is convenient in many situations including... video disc. In general, a search mode is convenient when the user interrupts the normal processing of the serial pipeline at a point where the machine does not expect such an...be the case.

In brief, the Huffman Decoder 321 works in conjunction with the other units shown in Figure 27. These other units are the Parser State Machine 322, the inshifter 323, the Index to Data unit 324, the ALU 325, and the Token Formatter 326. As described previously, connection between these blocks is governed by a two wire interface. A more detailed description of how these units function is subsequently described herein in greater detail, the focus here is on particular aspects... control certain functions of the Index to Data 324 and ALU 325. Control of these units by the Huffman Decoder is necessary for proper decoding of block-level information. Having the... further detail in the "More Detailed Description of the Invention" section.

The Index to Data unit 324 performs the second part of the multi-part algorithm. This unit contains a look up table that provides the actual Huffman decoded data. Entries in the...by detecting these in the Huffman Decoder 321, rather than in the Index to Data unit 324.

This index number is then passed to the Index to Data unit 324. In essence, the Index to Data unit is a look-up table. In accordance with one aspect of the algorithm, the lookup... ..format that JPEG specifies for transferring an alternate JPEG table.

From the Index to Data unit 324, the decoded index number or other data is passed, together with the accompanying control... ..the entering data to ensure that the DATA tokens are of the correct size for processing. In fact, the token stream can be corrected in some situations if the error is... ..an order that is useful for the decompression circuits, but not for the particular display unit being used. When a block of data enters the Buffer Manager, the Buffer Manager supplies...the output of the Spatial Decoder or Temporal Decoder and re-format it for a computer or display system. The details of this formatting will vary between applications. In a simple...Token. The DATA Token can have as many bits as are necessary for carrying out processing at a particular place in the system. All other Tokens ignore the extra bits.

A.3.2 The DATA Token

The DATA Token carries data from one processing stage to the next. Consequently, the characteristics of this Token change as it passes through... ..will be sufficient to collect DATA Tokens and to detect a few Tokens that provide synchronization information (such as PICTURE(underscore)START). In this regard, see subsequent sections A.16, "Connecting... ..from the data stream. This provides an alternative to doing the configuration via the micro processor interface.

A.3.4 Description of Tokens

This section documents the Tokens which are implemented... ..3.5.1. Note: JPEG requires a 2:1:1 structure for its macroblocks when processing 4:2:2 data. See Table A.3.5.

A.3.6 Special Token formats...either is low then the interface is taken to high impedance.

Note: on-chip data processing is not terminated when the DRAM interface is at high impedance. Therefore, errors will occur...decoded video's picture rate. Accordingly, this clock can be used to provide audio/video synchronization.

A.7.1 Spatial Decoder clock signals

The Spatial Decoder has two different (and potentially...in accordance with the present invention, must know what video standard is being input for processing. Thereafter, the system can accept either pre-existing Tokens or raw byte data which is...time a value is written into coded(underscore)data (7:0). Software is responsible for settling coded(underscore)extn to 0 before the last word of any Token is written to... ..0). The start of this new DATA Token then passes into the Spatial Decoder for processing.

Each time a new 8 bit value is written to coded(underscore)data (7:0... ..Detector analyses data in the DATA Tokens bit serially. The Detector's normal rate of processing is one bit per clock cycle (of coded(underscore)clock).

Accordingly, it will typically decode a byte of coded data every 8 cycles of coded(underscore)clock. However, extra processing cycles are occasionally required, e.g., when a non-DATA Token is supplied or when... ..Furthermore, this clock can be asynchronous to the main decoder(underscore)clock. Data transfer is synchronized to decoder(underscore)clock on-chip.

SECTION A.11 Start code detector

A.11.1...Code Detector. So, accessing these registers will be unreliable if the Start Code Detector is processing data. The user is responsible for ensuring that the Start Code Detector is halted before... ..Detector. In this case, the Tokens are passed through the Start Code Detector with no processing to other stages of the Spatial Decoder. These Tokens can only be inserted just before...result will be unpredictable if this is done when the Start Code Detector is actively

processing data.

Discard all mode can be safely initiated after any of the start Code Detector...start code non-alignment interrupt is suppressed.

In contrast, however, JPEG was designed for a computer environment where byte alignment is guaranteed. Therefore, marker codes should only be detected when byte... ..the other hand, was designed to meet the needs of both communications (bit serial) and computer (byte oriented) systems. Start codes in MPEG data should normally be byte aligned. However, the...result will be unpredictable if this is done when the Start Code Detector is actively processing data. So, before initiating a start code search, the start Code Detector should be stopped so no data is being processed. The Start Code Detector is always in this condition if any of the Start Code...the spatial video decoding circuits (inverse modeler, quantizer and DCT). This second logical buffer allows processing time to include a spread so as to accommodate processing pictures having varying amounts of data.

Both buffers are physically held in a single off... ..if the buffers are full or empty.

As stated in A.13.1.1, the unit for all the above mentioned registers is a 512 bit block of data. Accordingly, the...until there is space in the buffer. If a buffer continues to be full, more processing stages "up steam" of the buffer will halt until the Spatial Decoder is unable to... ..converting coded data into Tokens started by the Start Code Detector. There are four main processing blocks in the Video Demux: Parser State Machine, Huffman decoder (including an ITOD), Macroblock counter... ..or state machine follows the syntax of the coded video data and instructs the other units. The Huffman decoder converts variable length coded (VLC) data into integers. The Macroblock counter keeps...

Specification: ...values zero and 255 may not be used.

If such a picture were to be processed in a pipeline built in the practice of the present invention, then one of these...data must not be written over since it is data that must be saved for processing or use in a downstream device e.g., a pipeline stage, a device or a... ..connected to the pipeline upstream contains data D4 that is to be transferred into and processed in the pipeline. Stages B, D and E, in addition to the upstream device, contain... ..pipeline, in accordance with the preferred embodiments of the present invention, to "fill up" empty processing stages is highly advantageous since the processing stages in the pipeline ...data can be transferred into the pipeline and between stages even when one or more processing stages is blocked.

In the embodiment shown in Fig. 1, it is assumed that the... ..propagate all the way back to the beginning of the pipeline if there is some intermediate stage that is able to accept new data.

In the embodiment illustrated in Fig. 1...has been mentioned. It is to be further understood that each pipeline stage may also process the data it has received arbitrarily before passing it between its internal storage elements or... ..the portion of the pipeline that contains input and output storage elements and that arbitrarily processes data stored in its storage elements.

Furthermore, the "device" downstream from the pipeline Stage F... ..valid data, but also when a stage requires more than one dock phase to finish processing its data. This also can occur when it creates valid data in one or both... ..control the passage of data between adjacent storage elements. The VALID signal may also be processed in an analogous manner.

A great advantage of the two-wire interface (one wire for... ..In addition, two extra latches and a small number of gates are preferably added to process the ACCEPT and VALID signals that are associated with the data latches in each half... ..application so requires. The interface in accordance with this embodiment can also be used to process analog signals.

As discussed previously, while other conventional timing arrangements may be used, the interface...circuit B1, which may be provided to convert output data from input latch LDIN into intermediate data, which is then later loaded in an output data latch LDOUT, which comprises the... ..is connected either directly as an input to the validation output latch LVOUT, or via intermediate logic devices or circuits that may alter the signal.

Similarly, the output validation signal QVOUT... ..to the input of the validation input latch QVIN of the following stage, or via intermediate devices or logic circuits, which may alter the validation signal. This output QVIN is also...word.

Preferred Data Structure - "tokens"

In the sample application shown in Fig. 4, each stage processes all input data, since there is no control circuitry that excludes any stage from allowing...are connected together in a relatively simple configuration. The simplest configuration is a pipeline of processing steps. For example, in the one shown in Fig. 1. The use of tokens, however... ..flows from left to right in the diagram. Data enters the machine and passes into processing Stage A. This may or may not modify the data and it then passes the... ..advantage of the tokens is their ability to achieve this kind of communication. Since any processing stage that does not recognize a token simply passes it on unaltered to the next... ..is transmitted along with the address and data fields in each token so that a processing stage can pass on a token (which can be of arbitrary length) without having to... ..be the first word of a new token.

Note that although the simple pipeline of processing stages is particularly useful, it will be appreciated that tokens may be applied to more complicated configurations of processing elements. An example of a more complicated processing element is described below.

It is not necessary, in accordance with the present invention, to... ..has extension bits. An example of this is a token that activates a stage that processes video quantization values stored in a quantization table (typically a memory device). For example, a...turn, is of great importance in video data pipeline systems since it ensures that all processing stages can be continuously running at full bandwidth.

In accordance to the present invention, in...some other chips in the set. This is advantageous both from the perspective of a customer and from that of a chip manufacturer. Even if modifications mean that all chips are...the end of a token (and hence the start of the next token) to be processed correctly (Including simple non-manipulative transfer), even if the token is not recognized by the... ..block diagram of a pipeline stage whose function is as follows. If the stage is processing a predetermined token (known in this example as the DATA token), then it will duplicate... ..the address field of the DATA token. If, on the other hand, the stage is processing any other kind of token, it will delete every word. The overall effect is that...IN(underscore)VALID.

In the duplication stage, the output from the data latch LDIN forms intermediate data referred to as MID(underscore)DATA. This intermediate data word is loaded into the data output latch LDOUT only when an intermediate acceptance signal (labeled "MID(underscore)ACCEPT" in Fig. 8a) is set HIGH.

The portion of... ..data. These include a "DATA(underscore)TOKEN" signal that indicates that the circuitry is currently processing a valid DATA Token, and a NOT(underscore)DUPLICATE signal which is used to control duplication of data. When the circuitry is processing a DATA Token, the NOT(underscore)DUPLICATE signal toggles between a HIGH and a LOW... ..the token to be duplicated once (but no more times). When the circuitry is not processing a valid DATA Token then the NOT(underscore)DUPLICATE signal is held in a HIGH state. Accordingly, this means that the token words that are being processed are not duplicated.

As Fig. 8a illustrates, the upper six bits of 8-bit intermediate data word and the output signal QI1 from the latch LI1 form inputs to a... ..explained further below.

Latch LO1 performs the function of latching the last value of the intermediate extension bit (labeled "MID(underscore)EXTN" and as signal S4), and it loads this value... ..and the DATA(underscore)TOKEN signal will become "0", indicating that the circuitry is not processing a DATA token.

If QI1 is "0" and SO is "0", thereby indicating a DATA...phase and the DATA(underscore)TOKEN signal will become "1", indicating that the circuitry is processing a DATA token.

The NOT(underscore)DUPLICATE signal (the output signal QO3) is similarly loaded... ..LVOUT at the same time that MID(underscore)DATA is loaded into LDOUT and the intermediate extension bit (signal S4) is loaded into LEOUT. Signal S5 is also combined with the...above. This has the effect that all tokens except the one that causes the duplication process will be deleted from the token stream, since a device connected to the output terminals (OUTDATA, OUTEXTN and OUTVALID) will not recognize these token words as valid data.

As before... ..and is duplicated.

Referring now more particularly to Figure 10, there is shown a reconfigurable process stage.

Input latches 34 receive an input over a first bus 31. A first output from... ..the input latches 34 is passed as a first input over line 35 to a processing unit 36. A first output from the token decode subsystem 33 is passed over line 37 as a second input to the processing unit 36. A second output from the token decode 33 is passed over line 40 to an action identification unit 39. The action identification unit 39 also receives input from registers 43 and 44 over line 46. The registers 43... ..is determined by the history of tokens previously received. The output from the action identification unit 39 is passed over line 38 as a third input to the processing unit 36. The output from the processing unit 36 is passed to output latches 41. The output from the output latches 41 is Index to Data

Unit (ITOD) 64. The Huffman decoder 56 and the ITOD 64 work together as a single logical unit. The output from the ITOD 64 is passed over line 65 to an arithmetic logic

unit (ALU) 66. A first output from the ALU 66 is passed over line 67 to... ..blocks 133.

Referring to Figure 14b, in the JPEG and H.261 standards, the Common Intermediate Format (CIF) is used, wherein a picture 141 is encoded as 6 rows each containing... ..in a zigzag direction indicated by the arrow 144. The GOBs 142 are, in turn, processed row-by-row, left-to-right in each row.

Referring now to Figure 14c, it...in accordance with the practice of the present invention. A first picture 161 to be processed contains a first PICTURE(underscore)START token 162, first-picture information of indeterminate length 163, and a first PICTURE(underscore)END token 164. A second picture 165 to be processed contains a second PICTURE(underscore)START token 166, second picture information of indeterminate length 167... ..tokens 162 and 166 indicate the start of the pictures 161 and 165 to the processor. Likewise, the PICTURE(underscore)END tokens 164 and 168 signify the end of the pictures 161 and 165 to the processor. This allows the processor to process picture information 163 and 167 of variable lengths.

Referring to Figure 17, a split 171... ..Video Formatter (not shown in Figure 17).

Referring now to Figure 18, the prediction filtering process is illustrated. A forward picture 201 is passed over line 202 as a first input... ..the right of the value decode shift register 230, as indicated by area 231. This process eliminates overlapping start code images, as discussed below. A first output from the value decode...Code Detector. The Start Code Detector then receives a first data value image 244. Before processing the first data value image 244, the Start Code Detector may detect a second start... ..image 244 at a length 246. If this occurs, the Start Code Detector does not process the first data value image 244, and instead receives and processes a second data value image 247.

Referring now to Figure 22, a flag generator 251... line 1 of Table 600, whenever a "sequence start" image is received during H.261 processing or a "picture start" image is received during MPEG processing, the entire group of four control tokens is generated, each followed by its corresponding data...Picture Decoding

3. Motion Picture Decompression

4. RAM Memory Map

5. Bitstream Characteristics

6. Reconfigurable Processing Stage

7. Multi-Standard Coding

8. Multi-Standard Processing Circuit-2nd Mode of Operation

9. Start Code Detector

10. Tokens

11. DRAM Interface

12... described herein in greater detail) and reformatting this output for use, including display in a computer or other display systems, including a video display system. Implementation of this formatting varies significantly... the Spatial Decoder circuits.

The Spatial Decoder of the present invention performs all the required processing within a single picture. This reduces the redundancy within one picture.

The Temporal Decoder reduces... modeller 75, the inverse zig-zag 81 and the inverse DCT 83. The standard independent units within the Huffman decoder and parser include the ALU 66 and the token formatter 71.

Referring now to Figure 12, the standard-independent units include the DRAM interface 100, the fork 91, the FIFO register 96, the summer 98 and the output selector 106. The standard dependent units are the address generator 94, ...much of the operation is very similar between the three different compression standards.

The next unit is the state machine 68 (Figure 11) located within the Huffman decoder and parser. Here... The same holds true for JPEG, which is a third, completely independent program.

The next unit is the Huffman decoder 56 which functions with the index to data unit 64. Those two units cooperate together to perform the Huffman decoding. Here, the algorithm that is used for Huffman... to the Huffman decoder at different times consistent with the standard in operation.

The last unit on the chip that is dependent on the compression standard is the inverse quantizer 79...an H.261 group of blocks and an MPEG slice. When H.261 data is processed after the Start Code Detector, each group of blocks is preceded by a slice(underscore... these standards have totally different sets of tables.

As previously indicated, most of the system units are compression standard independent. If a unit is standard independent, and such units need not remember what CODING(underscore)STANDARD is being processed. All of the units that are standard dependent remember the compression standard as the CODING(underscore)STANDARD token flows... CODING(underscore)STANDARD tokens at the Start Code Detector that is positioned as the first unit in the pipeline, this change of compression standard is readily handled. The token says a... found in the standard, i.e. from the bitstream into a prediction mode token. This processing is performed by the Huffman decoder and parser state machine, where it is easy to... to that token. By having these tokens and using them appropriately, the design of other units in the machine is simplified. Although there may be some complications in the program, benefits...a first encoded signal (the MPEG or H.261 encoded video signal) in a pipeline processing system. The Temporal Decoder is not needed for JPEG decoding.

In this regard, the invention... the use of a single pipeline decoder and decompression system. The decoding and decompression pipeline processor is organized on a unique and special configuration which allows the handling of the multi... video signals through the use of techniques all compatible with the single pipeline decoder and processing system. The Spatial Decoder is combined with the Temporal Decoder, and the Video Formatter is... with only still pictures. The compression standard independent Spatial Decoder performs all of the data processing within the boundaries of a single picture. Such a decoder handles the spatial decompression of... to the multi-standard, configurable Video Formatter, which then provides an output to the display terminal. In a first sequence of similar pictures, each decompressed picture at the output of the... of control tokens and DATA tokens, in combination with a plurality of sequentially-positioned reconfigurable processing stages selected and organized to act as a standard-independent, reconfigurable-pipeline-processor.

With regard to JPEG decoding, a single Spatial Decoder with no off chip DRAM can... video. Accordingly, signals carried by DATA tokens pass directly through the Temporal Decoder without further processing when the Temporal Decoder is configured for a JPEG operation.

Another aspect of the present... for subsequent use in temporal decoding of subsequent pictures.

Generally, the Temporal Decoder performs the processing between pictures either-earlier and/or later in time with reference to the picture currently... is distributed among several areas of DRAM in the sense that the decompressed output information, processed by the Spatial Decoder, is stored in other DRAM registers by other random access memories... first decoder circuit (the Spatial Decoder) directly to the Video Formatter for handling without signal processing delay.

The Temporal Decoder also reorders the blocks of picture data for display by a... from a selection of pictures which have arrived earlier or later than the picture under processing. When a picture is described in this context, it may mean any one of the... 2. The result, i.e., the final decoded picture resulting from the addition of a process step performed by the decoder;

3. Previously decoded pictures read from the DRAM; and

4... START token and a subsequent PICTURE(underscore)END token.

After the picture data information is processed by the Temporal Decoder, it is either displayed or written back into a picture memory location. This information is then kept for further reference to be used in processing another different coded data picture.

Re-ordering of the MPEG encoded pictures for visual display... used to encode a referenced picture of a picture might be identified as being one unit long, another picture might be a number of units long, while still a third picture could be a fraction of that unit.

None of the existing standards (MPEG 1.2, JPEG, H.261) define a way of... picture rate, whereas the Video Formatter can handle a variable input picture rate.

6. RECONFIGURABLE PROCESSING STAGE

Referring again to Figure 10, the reconfigurable processing stage (RPS) comprises a token decode circuit 33 which is employed to receive the tokens... input latches 34. The output of the token decode circuit 33 is applied to a processing unit 36 over the two-wire interface 37 and an action identification circuit 39. The processing unit 36 is suitable for processing data under the control of the action identification circuit 39. After the processing is completed, the processing unit 36 connects such completed signals to the output, two-wire interface bus 40 through output... token decode circuit 33 are applied simultaneously to the action identification circuit 39 and the processing unit 36. The action identification function as well as the RPS is described in further detail... not standard independent circuits. The data flows through the token decode circuit 33, through the processing unit 36 and onto the two-wire interface circuit 42 through the output latches 41. If... wire interface 42 through the output

circuit 41. The present invention operates as a pipeline processor having a two-wire interface for controlling the movement of control tokens through the pipeline...time, the token decode circuit 33 provides a proper flag or index signal to the processing unit 36 to alert it to the presence of the token being handled by the action identification circuit 39.

Control tokens may also be processed.

A more detailed description of the various types of tokens usable in the present invention... standard now passing through the state machine shown with reference to Figure 10.

Similarly, the processing unit 36 which is under the control of the action identification circuit 39 is now ready to process the information contained in the data fields of the DATA token when it is appropriate... action identification circuit 39 and is immediately followed by a DATA token which is then processed by the processing unit 36. The control token exits the output latches circuit 41 over the output two-wire interface 42 Immediately preceding the DATA token which has been processed within the processing unit 36.

In the present invention, the action identification circuit, 39, is a state machine holding... show that the action can also be affected by the token that is currently being processed by the token decode circuit 33.

In general, there is shown token decoding and data processing in accordance with the present invention. The data processing is performed as configured by the action identification circuit 39. The action is affected by... information stored from previously decoded tokens in registers 43 and 44, the current token under processing, and the state and history information that the action identification unit 39 has itself acquired. A distinction is thereby shown between Control tokens and DATA tokens.

In any RPS, some tokens are viewed by that RPS unit as being Control tokens in that they affect the operation of the RPS presumably at... are viewed by the RPS as DATA tokens. Such DATA tokens contain information which is processed by the RPS in a way that is determined by the design of the particular... view of the same token. Some of the tokens might be viewed by one RPS unit as DATA Tokens while another RPS unit might decide that it is actually a Control Token. For example, the quantization table information... into a token called a quantization table token (QUANT(underscore)TABLE) which goes down the processing pipeline. As far as that machine is concerned, all of that was data; it was... sort of data into another sort of data, which is clearly a function of the processing performed by that portion of the machine. However, when that information gets to the inverse... present. This information is viewed as control information, and then that control information affects the processing that is done on subsequent DATA tokens because it affects the number that you multiply... important feature of the invention is that each of the stages of circuitry has the processing capability within it to be able to perform the necessary operations for each of the... operations are to be performed at a given time, come as tokens. There is one processing element that differs between the different stages to provide this capability. In the state machine... standard is and it looks up the parameters that it needs to apply to the processing elements in order to perform a proper operation. For example, the inverse quantizer will look... is set to 1 for a particular compression standard, and will apply that to its processing circuitry.

In a similar sense the Huffman decoder 56 has a number of tables within... MPEG video standard or the JPEG video standard. These three compression coding standards specify similar processes to be done on the arriving data, but the structure of the datastreams is different... token stream embodying the current coding standard. The control tokens are passed through the pipeline processor, and are used, i.e., decoded, in the state machines to which they are relevant... this regard, the DATA Tokens are treated in the same fashion, insofar as they are processed only in the state machines that are configurable by the control tokens into processing such DATA Tokens. In the remaining state machines, they pass through unchanged.

More specifically, a... signals. The remaining portions of the token are used to indicate and identify the internal processing control function which is standard for all of the datastreams passing through the pipeline processor. In one form of the invention, the token extension is used to carry the current... accompanying data. As previously discussed, this information is utilized in the system to reconfigure the processing stage used to perform the function required by the various standards created for that purpose... picture number as indicated by the value.

The system also includes a multi-stage parallel processing pipeline operating under the principles of the two-wire interface previously described. Each of the... the token presently entering the state machine into the action identification circuit 39 or the processing unit 36, as appropriate. The processing unit has been previously reconfigured by the next previous control token into the form needed for handling the current coding standard, which is now entering the processing stage and carried by the next DATA token. Further, in accordance with this aspect of the invention, the succeeding state machines in the processing pipeline can be functioning under one coding standard, i.e., H.261, while a previous... tokens required to decode a number of coding standards with a fixed number of reconfigurable processing stages. More specifically, the PICTURE(underscore)END control token is employed because it is important... standard machine, it is necessary to create additional control tokens within the multi-standard pipeline processing machine which will then indicate which one of the standard decoding techniques to use. Such... and to push the current picture through the decoder to the display.

8. MULTI-STANDARD PROCESSING CIRCUIT - SECOND MODE OF OPERATION

A compression standard-dependent circuit, in the form of the... of the Start Code Detector will subsequently be discussed in further detail, as will the process of starting up of the decoder.

The aforementioned description has been concerned primarily with the... the data which immediately follows according to the standard. However, in the multi-standard pipeline processing system of the present invention, where compatibility is required for multiple standards, the system has... signals, including flag signals, are generated by each state machine to handle some of the processing within that state machine. Values carried in the standards can be used to access machine... its contents must be removed from the two wire interface to ensure that no further processing takes place using these 3 bytes. The decode register is emptied, and the value decode... 10. TOKENS

In the practice of the present invention, a token is a universal adaptation unit in the form of an interactive interfacing messenger package for control and/or data functions and is adapted for use with a reconfigurable processing stage (RPS) which is a stage, which in response to a recognized token, reconfigures itself to perform various operations.

Tokens may be either position dependent or position independent upon the processing stages for performance of various functions. Tokens may also be metamorphic in that they can be altered by a processing stage and then passed down the pipeline for performance of further functions. Tokens may interact... other functions, and the specific interaction with a stage may be conditioned by the previous processing history of a stage.

A PICTURE(underscore)END token is a way of signalling the... through a fixed size, fixed width buffer.

The present invention is directed to a pipeline processing system which has a variable configuration which uses tokens and a two-wire system. The... do not use control tokens.

The control tokens are generated by circuitry within the decoder processor and emulate the operation of a number of different type standard-dependent signals passing into the serial pipeline processor for handling. The technique used is to study all the parameters of the multi-standards that are selected for processing by the serial processor and noting 1) their similarities, 2) their dissimilarities, 3) their needs and requirements and 4) selecting the correct token function to effectively process all of the standard signals sent into the serial processor. The functions of the tokens are to emulate the standards. A control token function is... the standard dependent signals and as an element to transmit control information through the pipeline processor.

In prior art system, a dedicated machine is designed according to well-known techniques to... tokens provide and make a sensible format for communicating information through the decompression circuit pipeline processor. In the design selected hereinafter and used in the preferred embodiment, each word of a... However, this is not a limitation on the invention, but on the magnitude of the processing steps elected to be accomplished by use of these tokens. It is to be noted... bit address for use in accessing the random access memories used throughout this serial decompression processor. This provides an additional degree of variability that facilitates a broad range of versatility.

As previously described, the DATA token carries data from one processing stage to the next. Consequently, the characteristics of this token change as it passes through... longest number of data bits because it needs to provide the most information to the processing unit so that it can start the decompression with as much information as possible. Words which... to receive an address, it waits for the address generator to supply a valid address, processes that address and then sets the accept line high for one dock period. Thus, it... be read. This signal passes between two asynchronous dock regimes and, therefore, passes through three synchronizing flip flops.

Provided RAM2 312 is empty, the next item of data to arrive on... interesting.

In general, prediction data will be offset from the position of the block being processed as specified in the motion vectors in x and y. Thus, the block of data... address, 9. Data is read from this address and the x value is incremented. The

process is repeated until the x value reaches its stop value, at which point, the y... data is read, the value is again incremented until it reaches its stop value. The process is repeated until both x and y values have reached their stop values. Thus, the... invention, is that additional information must be provided to the prediction filters to indicate what processing is required on the data. This consists of the following:

a "last byte" signal indicating... bit 0) is incremented and the x address (3 LSBS) is reset to zero. This process is repeated until 64 bytes have been read. With a 16 or 32 bit wide... register while its access register is set to zero, the results are undefined.

14. MICRO-PROCESSOR INTERFACE

A standard byte wide micro-processor interface (MPI) is used on all circuits with in the Spatial Decoder and Temporal Decoder... the parameter column. The actual specifications are shown in the respective columns min, max and units.

The DC operating conditions can be seen with reference to Table A.6.3. Here... the signal is present the maximum amount of time that this signal is available. The Units column gives the units of measurement used to describe the signals.

16. MPI WRITE TIMING

The general description of... Consequently, the machine will not go into error recovery mode and will successfully continue to process the coded data.

A still further advantage of the use of a PICTURE(underscore)END token is that the serial pipeline processor will continue the processing of uninterrupted data. Through the use of a PICTURE(underscore)END token, the serial pipeline processor is configured to handle less than the expected amount of data and, therefore, continues processing. Typically, a prior art machine would stop itself because of an error condition. As previously... of the Huffman decode and Video Demultiplexor know the number of blocks that it will process during each picture recovery cycle. When the correct number of blocks do not arrive from... Each of the state machines recognizes a FLUSH control token as information not to be processed. Accordingly, the FLUSH token is used to fill up all of the

remaining empty parts... less information than normally expected to decode the last picture. The Huffman decode circuit finishes processing the information contained in the last picture, and outputs this information through the DRAM interface... token, in accordance with the present invention, is used to pass through the entire pipeline processor and to ensure that the buffers are emptied and that other circuits are reconfigured to... underscore)END token, a padding word and a FLUSH token indicating to the serial pipeline processor that the picture processing for the current picture form is completed. Thereafter, the various state machines need reconfiguring to... FLUSH token resets each stage as it passes through, but-allows subsequent stages to continue processing. This prevents a loss of data. In other words, the FLUSH token is a variable... AFTER PICTURE

The STOP(underscore)AFTER(underscore)PICTURE function is employed to shut down the processing of the serial pipeline decompressing circuit at a logical point in its operation. At this... a picture, the STOP(underscore)AFTER(underscore)PICTURE operation signals the end of all current processing.

22. MULTI(underscore)STANDARD - SEARCH MODE

Another feature of the present invention is the use... underscore)MODE control token which is used to reconfigure the input to the serial pipeline processor to look at the incoming bit stream. When the search mode is set, the Start... combination of control tokens, and DATA tokens along with the reconfiguration circuits, to provide similar processing.

The use of search mode in the present invention is convenient in many situations including... video disc. In general, a search mode is convenient when the user interrupts the normal processing of the serial pipeline at a point where the machine does not expect such an... be the case.

In brief, the Huffman Decoder 321 works in conjunction with the other units shown in Figure 27. These other units are the Parser State Machine 322, the inshifter 323, the Index to Data unit 324, the ALU 325, and the Token Formatter 326. As described previously, connection between these blocks is governed by a two wire interface. A more detailed description of how these units function is subsequently described herein in greater detail, the focus here is on particular aspects... control certain functions of the Index to Data 324 and ALU 325. Control of these units by the Huffman Decoder is necessary for proper decoding of block-level information. Having the be used by the Huffman Decoder for all three standards.

The Index to Data unit 324 performs the second part of the multi-part algorithm. This unit contains a look up table that provides the actual Huffman decoded data. Entries in the... by detecting these in the Huffman Decoder 321, rather than in the Index to Data unit 324.

This index number is then passed to the Index to Data unit 324. In essence, the Index to Data unit is a look-up table. In accordance with one aspect of the algorithm, the look... format that JPEG specifies for transferring an alternate JPEG table.

From the Index to Data unit 324, the decoded index number or other data is passed, together with the accompanying control... the entering data to ensure that the DATA tokens are of the correct size for processing. In fact, the token stream can be corrected in some situations if the error is... an order that is useful for the decompression circuits, but not for the particular display unit being used. When a block of data enters the Buffer Manager, the Buffer Manager supplies... the output of the Spatial Decoder or Temporal Decoder and re-format it for a computer or display system. The details of this formatting will vary between applications. In a simple... Token. The DATA Token can have as many bits as are necessary for carrying out processing at a particular place in the system. All other Tokens ignore the extra bits.

A.3.2 The DATA Token

The DATA Token carries data from one processing stage to the next. Consequently, the characteristics of this Token change as it passes through... will be sufficient to collect DATA Tokens and to detect a few Tokens that provide

synchronization information (such as PICTURE(underscore)START). In this regard, see subsequent sections A.16, "Connecting... ..from the data stream. This provides an alternative to doing the configuration via the micro processor interface.

A.3.4 Description of Tokens

This section documents the Tokens which are Implemented... ..3.5.1. Note: JPEG requires a 2:1:1 structure for its macroblocks when processing 4:2:2 data. See Table A.3.5.

A.3.6 Special Token formats...either is low then the interface is taken to high impedance.

Note: on-chip data processing is not terminated when the DRAM interface is at high impedance. Therefore, errors will occur...decoded video's picture rate. Accordingly, this clock can be used to provide audio/video synchronization.

A.7.1 Spatial Decoder clock signals

The Spatial Decoder has two different (and potentially...in accordance with the present invention, must know what video standard is being input for processing. Thereafter, the system can accept either pre-existing Tokens or raw byte data which is... ..time a value is written into coded(underscore)data (7:0). Software is responsible for settling coded(underscore)extn to 0 before the last word of any Token is written to... ..0). The start of this new DATA Token then passes into the Spatial Decoder for processing.

Each time a new 8 bit value is written to coded(underscore)data (7:0...Detector analyses data in the DATA Tokens bit serially. The Detector's normal rate of processing is one bit per clock cycle (of coded(underscore)clock). Accordingly, it will typically decode a byte of coded data every 8 cycles of coded(underscore)clock. However, extra processing cycles are occasionally required, e.g., when a non-DATA Token is supplied or when... ..Furthermore, this clock can be asynchronous to the main decoder(underscore)clock. Data transfer is synchronized to decoder(underscore)clock on-chip.

SECTION A.11 Start code detector

A.11.1... ..Code Detector. So, accessing these registers will be unreliable if the Start Code Detector is processing data. The user is responsible for ensuring that the Start Code Detector is halted before... ..Detector. In this case, the Tokens are passed through the Start Code Detector with no processing to other stages of the Spatial Decoder. These Tokens can only be inserted just before...result will be unpredictable if this is done when the Start Code Detector is actively processing data.

Discard all mode can be safely initiated after any of the Start Code Detector... ..start code non-alignment interrupt is suppressed.

In contrast, however, JPEG was designed for a computer environment where byte alignment is guaranteed. Therefore, marker codes should only be detected when byte... ..the other hand, was designed to meet the needs of both communications (bit serial) and computer (byte oriented) systems. Start codes in MPEG data should normally be byte aligned. However, the...result will be unpredictable if this is done when the Start Code Detector is actively processing data. So, before initiating a start code search, the Start Code Detector should be stopped so no data is being processed. The Start Code Detector is always in this condition if any of the Start Code...will only occur if very short streams are being decoded or if the off-chip buffers are very large as compared to the picture format being decoded).

In Figure 69 stream... ..the spatial video decoding circuits (inverse modeler, quantizer and DCT). This second logical buffer allows processing time to include a spread so as. to accommodate processing pictures having varying amounts of data.

Both buffers are physically held in a single off...if the buffers are full or empty.

As stated in A.13.1.1, the unit for all the above mentioned registers is a 512 bit block of data. Accordingly, the...
...until there is space in the buffer. If a buffer continues to be full, more processing stages "up steam" of the buffer will halt until the Spatial Decoder is unable to... ..converting coded data into Tokens started by the Start Code Detector. There are four main processing blocks in the Video Demux: Parser State Machine, Huffman decoder (including an ITOD), Macroblock counter... ..or state machine follows the syntax of the coded video data and instructs the other units. The Huffman decoder converts variable length coded (VLC) data into integers. The Macroblock counter keeps...In the present invention, picture dimensions are described to the Spatial Decoder in 2 different units: pixels and macroblocks. JPEG and MPEG both communicate picture dimensions in pixels. Communicating the dimensions... ..v and max(underscore)component(underscore)id specify the composition of the macroblocks (minimum coding units in JPEG). Each is a 2 bit register than can hold values in the range...supports some picture formats beyond those defined by JPEG and MPEG.

JPEG limits minimum coding units so that they contain no more than 10 blocks per scan. This limit does not apply to the Spatial Decoder since it can process any minimum coding unit that can be described by blocks(underscore)h(underscore)n, blocks(underscore)v(underscore)n... ..for 4:2:0 macroblocks (see Table A.14.8). However, the Spatial Decoder can process three other component macroblock structures, (e.g., 4:2:2.

A.14.5 Video events...of the Token buffer and the output of the Spatial Decoder.

There are three main units responsible for spatial decoding: the inverse modeler, the inverse quantizer and the inverse discrete cosine... ..At this point, the values in the DATA Tokens are quantized coefficients.

The inverse modelling process is the same regardless of the coding standard currently being used. No configuration is required...

Claims: ...system having an inverse modeller stage and an inverse discrete cosine transform stage, comprising a processing stage, positioned between said inverse modeller stage and said inverse discrete cosine transform stage, responsive to tokens for processing data, wherein said tokens each comprise a plurality of data words, each said word including... ..can be unlimited; wherein said tokens are communicated from said inverse modeller stage to said processing stage.

2. The system as recited in claim 1, wherein said token is a QUANT... ..9. The system as recited in any one of claims 1 to 8, wherein said processing stage is connected to said inverse modeller stage and said inverse discrete cosine transform stage...

6/K/15 (Item 15 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Country	Number	Kind	Date		
Type	Pub. Date		Kind	Text	
Available Text	Language		Update	Word Count	
Total Word Count (Document A)					

Total Word Count (Document B)
Total Word Count (All Documents)

Specification: ...and more particularly to a system and method which provides for broadband information communication between processor-based systems through a centralized communication array. BACKGROUND OF THE INVENTION

In the past, information communication between processor-based systems, such as local area networks (LAN) and other general purpose computers, separated by significant physical distances has been an obstacle to integration of such systems. The... ..lack the fault tolerance or reliability found in systems designed for reliable transmission of important processor-based system information.

Another historically available group of communication choices is found at the opposite... ..and therefore do not lend themselves to simple interfacing with a variety of general purpose processor-based systems.

Although a fibre optic ring provides economy if utilized by a plurality of systems, it must be physically coupled to such systems. As the cost of purchasing, placing, and maintaining such a ring is, great, even the economy of multi-system utilization... ..information communication for a communication system providing cost effective bridging of large physical distances between processor-based systems.

A further need exists in the art for a communication system providing high speed broadband information communication between processor-based systems.

A still further need exists in the art for a fault tolerant communication system providing reliable bridging of physical gaps between processor-based systems.

Additionally, a need exists in the art for a broadband communication system providing simple connectivity to a variety of processor-based systems and communication protocols, including general purpose computer systems and their standard communication protocols.

SUMMARY OF THE INVENTION

These and other objects, needs... ..communication array, or hub, is centrally located to provide an air link between physically separated processor-based systems, or other sources of communication such as voice communication, utilizing a communication device... ..comprises a plurality of individual antenna elements in time division multiplex (TDM) communication with a processor-based system. This system processes signals received at each antenna element in order to route them to their desired destination... ..taken in conjunction with the accompanying drawings, in which:

FIGURE 1 illustrates the interconnection of processor-based systems of a preferred embodiment of the present invention;

FIGURE 2A illustrates an isometric... ..centralized communication array and nodes of the present invention;

FIGURE 6 illustrates the interconnection of processor-based systems through a network of hubs of the present invention; and

FIGURES 7-8... ..for example, to provide high speed bridging of a physical gap between a plurality of processor-based systems, as illustrated by system 100. The processor-based systems may include local area networks (LAN), such as LANs 110 and 120, or individual computer systems, such as PC 130. It shall be appreciated that the processor-based systems utilizing the present invention may be general purpose computers, both standing alone and interconnected such as by a LAN. Furthermore, the system can connect... ..or video in combination with, or in place of, communication sourced by the above mentioned processor-based systems.

between the various signals modulated... communication. Therefore, each antenna module 220 may include TDD switches 890 and 891 coupled to synchronizer 830 to provide synchronous switching the antenna element during transmit and receive frames, as is... and/or down-conversion of the signal. For example, in the preferred embodiment where an intermediate frequency of 400-500 MHz and a radio frequency of approximately 38 GHz are used... to up-convert and/or down-convert the signal in stages, such as through an intermediate frequency of 3 GHz. Therefore, in the preferred embodiment, converters 892 and 893 include multiple... between 400-500 MHz, 3 GHz, and 38 GHz.

It shall be understood that an intermediate frequency closer to the radio frequency may be utilized, thus eliminating the need for both... matrix suitable for lower frequencies than for higher frequencies. Therefore, in the preferred embodiment, an intermediate frequency significantly lower than the radio frequency to be transmitted is utilized.

In the preferred... may be utilized according to the present invention.

In addition to communication of information between processor-based systems through hub 101, control functions may also be communicated between hub 101 and... 101 through a backbone, such as backbone 160 illustrated in FIGURE 6, ultimately to other

processor-based systems. It shall be understood that a plurality of such backbone communications means may... antenna element, when switched in communication with controller 250, ultimately to be received by another processor-based system. Directing attention again to FIGURE 6, this communication path is illustrated, for example... network 110 communicating through hub 101 to network 120.

Larger geographical distances between two communicating processor-based systems may be bridged by utilization of multiple hubs. For example, as illustrated in... link via antenna elements. These two hubs may provide information communication between any combination of processor-based systems in communication with either hub.

It shall be appreciated that information received by... period, or channel of hub 101. Such an embodiment is efficient where, for example, a processor-based system, in communication with hub 101 through antenna element 200a, is only desirous of communicating with a processor-based system, in communication with hub 101 through element 200b.

However, where a processor-based system is desirous of communicating through hub 101 with a plurality of different processor-based systems, or a single antenna element is utilized by a plurality of processor-based systems, the above described correlation table may be ineffective. Therefore, in a preferred embodiment... fully illustrated. In a preferred embodiment node 150 is comprised of two primary components, outdoor unit 410 and indoor unit 450, as depicted in FIGURE 4.

Outdoor unit 410 includes antenna 420, module 430 and modem 440. Where EHF is used, antenna 420... the link illustrated between CPU 460 and module 430 may provide a signal controlling the synchronized switching the synchronized switching of the TDD switches according to a TDD frame of an associated hub. Modem... stated above where, for example, a different carrier frequency or beam pattern is desired.

Indoor unit 450 includes CPU 460, RAM 470 and interface 480. It shall be understood that indoor unit 450 and outdoor unit 410 are coupled such that information received by antenna 420 as RF energy is communicated to indoor unit 450.

Interface 480 provides data communication between indoor unit 450, and thus node 150, and a processor-based system such as LAN 490 illustrated in FIGURE 4. Furthermore, interface 480 formats the data communication to be compatible with the processor-based system so coupled. As for example, where LAN 490 is coupled to node 150... 490 utilizes Ethernet compatible communication protocol. However, where node 150 is coupled to a single

computer, it may be advantageous for interface 480 to provide asynchronous receive/transmit protocol. It shall...
 ...260 or 460, or may be integral to the particular transmission protocol utilized by the processor based systems as,
 for example, data packets conforming to Ethernet protocol. Regardless of its source...tuner within antenna module
 430. Such a control signal may be provided by the control processor to program phase lock loop circuitry, or
 synthesizer hardware, within the antenna module to select... ..of communication available for communication
 between node 150 and hub 101 due to TDM, and synchronizing information, such as frame timing and propagation
 delay offset, to enable TDM and/or TDD... ..above mentioned communication instructions. Of course, such an
 initialization algorithm may be stored in a processor-based system in communication with node 150 to achieve the
 same results if desired.

The... ..node 150, the initialization algorithm utilized by hub 101 alternatively may be stored in a processor-based
 system in communication with hub 101 to achieve the same results. The initialization algorithm... ..in RAM 470 to
 enable CPU 460 to time transmission through antenna 410 to achieve synchronization with the switching of antenna
 elements by ODU controller 230. Of course, it may not... ..delay, and therefore the distance between node and hub,
 is determined by the node initially synchronizing to the frame timing established by the hub. Thereafter, the node
 transmits a shortened burst... ..channels available for communication between a specific node as well as timing
 information to allow synchronization of communication between the node and the TDM antenna element of hub
 101.

The timing... ..minimize the potential for co-channel interference and, to a certain extent, multi-path interference,
 synchronization of transmission and reception at each antenna element is desirable. For example each antenna
 element... ..a predetermined Rx frame. Likewise, each hub of a network of such hubs may be synchronized to
 transmit and receive only during the same predetermined Tx and Rx frames. It shall...

Claims: ...via at least a second frequency in the millimeter wave spectrum of frequencies; and a processor-based
 communication hub (101, 610, 630) comprising:a plurality of hub antennas (200), each hub...

Claims: ...le spectre de frequences des ondes millimetriques ; et un concentrateur de communication a base de
 processeur (101, 610, 630) comprenant :une pluralite d'antennes de concentrateur (200), chaque antenne de
 concentrateur...

6/K/16 (Item 16 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

System fur die Synchronisation und Integritat von Passworter in einer DCE Umgebung

Country	Number	Kind	Date		
Type		Pub. Date		Kind	Text
Available Text		Language		Update	Word Count
Total Word Count (Document A)					
Total Word Count (Document B)					
Total Word Count (All Documents)					

Specification:

The present invention relates generally to network system servers and, more particularly, to password maintenance and resource sharing among data processing systems with password synchronization. More specifically still, the present invention relates to providing password synchronization and integrity between a general repository and a plurality of users within a client/server system.

Computer networks are well known in the arts and continue to grow in both size and complexity. This growth is fuel led by more computers being connected to networks and connecting networks to other networks. This is done in order to enable computers to work together efficiently, to simplify sharing resources such as files, applications, printers, and even special computers.

Unfortunately, many networks contain computers from different manufacturers, which complicates the task of getting computers to work together efficiently. Computers in such "multi-vendor" networks are usually difficult to operate together because they do not... ..mechanisms. The lack of a common network naming scheme also limits the degree to which computers can share information.

In order to overcome many of the above problems, the Open System... ..formed that has developed the Distributed Computing Environment ("DCE") that transforms a group of networked computers into a single, coherent computing engine. DCE masks differences among different kinds of computers, thereby enabling users to develop and execute distributed applications that can tap into a network... ..parts of these applications can execute concurrently, they can be much more powerful than single processor applications that must act on data sequentially.

Unfortunately, even distributed computing environments have their own... ..protect data that must be shared among multiple users for one. Additionally, events must be synchronized between separate computers and computers with differing data formats and file-naming schemes must be allowed to work cooperatively one... ..along with a distributed set of daemons and libraries, compose the DCE security service.

When servers enforce security, each client must provide its user's identity and access rights. These are... ..access to every DCE resource - directories, files, printers, and so on - is controlled by a server, the server's demands for authentication and authorization require comprehensive network security. This applies to DCE servers, such as those of the CDS, as well as user application servers. The server verifies the clients' authenticity and authorization before allowing access to the resource.

Typically, security service... ..means that the clients call local security routines to request authentication information from a security server and pass it to other servers. Servers also call security routines to verify authentication information and enforce authorization. Authentication is the ability... ..able to propagate the plain text password securely from the DCE registry to registries that synchronize with the DCE registry so that other registries can then encrypt the password in their... ..has occurred.

In accordance with the present invention, there is now provided a network system server for providing password composition checking for a plurality of clients, the server comprising: a main data store; a security server, coupled to said main data store and said plurality of clients; a password synchronization server, coupled to said security server; a plurality of password strength servers, coupled to said password synchronization server, that provides password integrity among said plurality of clients so that each client maintains a password whose composition is consistent with said network server system.

Viewing the present invention from another aspect, there is now provided a network system server that provides password synchronization between a main data store and a plurality of secondary data stores, comprising: a security server, coupled to said main data store; a plurality of clients, coupled to said security server for accessing said main data store wherein each client maintains a unique, modifiable password; a password synchronization server, coupled to said security server and to said plurality of secondary data stores; and a password repository, coupled to said password synchronization server, that stores said passwords whereby one of said secondary data stores can retrieve said passwords via said password synchronization server so that each client is able to maintain a single,

unique password among said plurality... ..Viewing the present invention from yet another aspect, there is now provided a network system server for providing password synchronization between a main data store and a plurality of secondary data stores, the server comprising: a security server, coupled to said main data store; a plurality of clients, coupled to said security server for accessing said main data store wherein each client maintains a unique, modifiable password; and a password synchronization server, coupled to said security server and to said plurality of secondary data stores, that provides password propagation synchronization to each of said secondary data stores from a user associated with one of said... ..clients so that each client maintains a password whose composition is consistent with a network server system in which said plurality of clients operate.

Viewing the present invention from another aspect, there is now provided a method for providing password synchronization between a main data store and a plurality of secondary data stores, comprising: storing in... ..password selected by a user associated with one of a plurality of clients; propagating password synchronization to each of said secondary data stores from said main data store so that said... ..present invention from yet another aspect, there is now provided a method for providing password synchronization between a main data store and a plurality of secondary data stores, comprising: maintaining a... ..stores.

In a preferred embodiment of the present invention, there is provided a network system server that provides password composition checking for a plurality of clients is disclosed. The network system server includes a main data store, a security server, which is coupled to the main data store and the plurality of clients, a password synchronization server, which is coupled to the security server, a plurality of password strength servers, each of which is coupled to the password synchronization server, that provides password integrity among the plurality of clients so that each client maintains a password whose composition is consistent with the network server system. Each of the password strength servers is uniquely programmable with respect to performing password composition checking.

In another embodiment of the present invention, there is provided a network system server that provides password synchronization between a main or global registry or main data store and a plurality of foreign or secondary registries or secondary data stores. The network server further includes a security server, which is coupled to the main data store, a plurality of clients, coupled to the security server for accessing the main data store wherein each client maintains a unique, modifiable password, and a password synchronization server, coupled to the security server and the plurality of secondary data stores, that provides password propagation synchronization to each of the secondary data stores from a user associated with one of the...computing environment and each of the plurality of secondary data stores is a foreign registry server having an associated permanent secondary permanent database or long term memory storage. The password propagation... ..current password status of the secondary data stores.

The present invention thus provides a password synchronization function that allows foreign registries to maintain local password synchronization with those maintained by a security server in the main registry. The function allows any number of foreign registries to automatically receive... ..system services in order to receive such password updates that allow them to maintain password synchronization. The main registry acts as a central repository for the information controlling which password changes... ..plurality of clients that includes a user's password, an information account for the password synchronization server that binds each of the plurality of secondary data stores with each of the plurality of clients, and an information account for each of the secondary data stores.

The security server provides a clear-text password to the password synchronization server for propagating to each of the plurality of secondary data stores. Further, the security server includes means for encrypting the clear-text password for storage in the information account.

A... ..memory and retry data store retry queue is further provided that couples to the password synchronization server. The temporary memory and retry data store contains information supporting a propagation retry that allows

the password synchronization server to perform a propagation retry in the event of a temporary foreign registry or password synchronization server outage.

In yet another embodiment of the present invention, there is provided a network system server that provides password synchronization between a main data store and a plurality of secondary data stores is disclosed. The network system server includes a security server, which is coupled to the main data store, a plurality of clients, which are coupled to the security server for accessing the main data store wherein each client maintains a unique, modifiable password, a password synchronization

server, which is coupled to security

server and the plurality of secondary data stores, and a password repository, which is coupled to the password synchronization server, that stores the passwords. Any one of the secondary data stores can retrieve the passwords via the password synchronization server so that each client is able to maintain a single, unique password among the plurality... ..accompanying drawings, wherein:

Figure 1 is a block diagram of a prior art DCE security server;

Figure 2 depicts a pictorial representation of a distributed data processing system according to the present invention;

Figure 3 is an example of a DCE synchronization and security system for use in the distributed processing system of Figure 2;

Figure 4 illustrates the password synchronization retrieve function;

Figure 5 illustrates the password synchronization propagation function;

Figure 6 depicts the steps that extend password synchronization to support password strength servers that may be customer-tailored according to a user's particular needs;

Figure 7 depicts a flow chart of user password processing from the security server;

Figure 8 illustrates the password synchronization server propagation queue thread;

Figure 9 depicts a flow chart of the password synchronization server retry queue thread according to the present invention.

With reference now to the figures, and... ..with reference to Figure 2, there is depicted a pictorial representation of a distributed data processing system 8 or distributed computing environment (DCE) system, which may be utilized to implement the method and system of the present invention. As may be seen, distributed data processing system 8 may include a plurality of local resource networks, such as Local Area Networks (LAN) 10 and 32, each of which preferably includes a plurality of individual computers 12 and 30, respectively. Of course, those skilled in the art will appreciate that a plurality of Intelligent Work Stations (IWS) coupled to a host processor may be utilized for each such network.

As is common in such data processing systems, each individual computer may be coupled to a storage device 14 and/or a printer/output device 16... ..invention, to store the various data objects or documents which may be periodically accessed and processed by a user within distributed data processing system 8, in accordance with the method and system of the present invention. In a manner well known in the prior art, each such data processing procedure or document may be stored within a storage device 14 which is associated with... ..objects associated therewith.

Still referring to Figure 2, it may be seen that distributed data processing system 8 may also include multiple mainframe computers, such as mainframe computer 18, which may be preferably coupled to Local Area Network

(LAN) 10 by means of communications link 22. Mainframe computer 18 may also be coupled to a storage device 20 which may serve as remote... ..Area Network (LAN) 10 via communications controller 26 and communications link 34 to a gateway server 28. Gateway server 28 is preferably an individual computer or Intelligent Work ...Local Area Network (LAN) 32 and Local Area Network (LAN) 10, a plurality of data processing procedures or documents may be stored within storage device 20 and controlled by mainframe computer 18, as Resource Manager or Library Service for the data processing procedures and documents thus stored.

Of course, those skilled in the art will appreciate that mainframe computer 18 may be located a great geographical distance from Local Area Network (LAN) 10 and... ..in California while Local Area Network (LAN) 10 may be located within Texas and mainframe computer 18 may be located in New York.

As will be appreciated upon reference to the foregoing, it is often desirable for users within one portion of distributed data processing network 8 to access a data object or document stored in another portion of data processing network 8. In order to maintain a semblance of order within the documents stored within data processing network 8 it is often desirable to implement a one-way access control program. This... ..to a document within a Resource Manager or Library Service. In this manner, the data processing procedures and documents may be accessed by enrolled users within distributed data processing system 8 and periodically "locked" to prevent access by other users. It is the system... ..only, but is extendible to all distributed LANs that use a main or global system server that can accommodate foreign or remote registry servers. Accordingly, although DCE is used in its properly accepted definition, that definition is extended in this disclosure to mean extended LANs or WANs having a plurality of non-common remote servers attached to a main system server for allowing a user or plurality of users to maintain password synchronization across the non-common remote servers from their particular local host server.

The DCE system of Figure 2 includes a mechanism for an enterprise to install a password strength checking server. Whenever the DCE registry receives a request to update a password, which includes the user... ..record, which includes the user name and the proposed new password, to the password strength server. The data is then forwarded in such a way that it is encrypted over the... ..the DCE's RPC, but the plain text password can be decrypted by the strength server. The strength server typically checks the password against the rules determined by the customer to insure that the password is not easily guessed. Such rules are well known in the arts and are left to the skilled artisan to implement. The password strength server then replies back to the registry service, informing it that the change is either acceptable or not acceptable. An example of a DCE synchronization and security system for use in the distributed processing system of Figure 2 is illustrated in Figure 3.

For propagation of plain text passwords to other security registries that wish to be synchronized with DCE, a novel password synchronization server is now provided that securely propagates plain text passwords to these registries. A model of the DCE system with appropriate DCE registry and the novel password synchronization server is illustrated in the block diagram of Figure 3. In Figure 3, DCE registry 102 is coupled to a DCE security server 104, which in turn is coupled to a password synchronization server 106 and various foreign registry servers X-Z 108. Each foreign registry server 108 typically comprises a local server and a secondary data store used to maintain local information such as local passwords for one.

Password synchronization server 106 is further coupled to various local password strength servers X-Z 110 as well as to a password repository 112, which includes user names and associated passwords. DCE security server 104 is further coupled to each client W-Y 114.

Within DCE registry 102 there... ..each client Y are various ERAs such as password management binding, which allows the security server to locate the password synchronization server or password strength server for this client; foreign registry, which enumerates the foreign registries authorized to receive password propagations... ..allows a foreign registry to locate and request the plain text password from the password synchronization server; and password strength, which identifies the password strength server for this client. Additionally, accounts for each foreign registry 108 are also maintained within DCE... ..registry Z, all foreign registries from X-Z, or any number

desired by a particular customer, are provided in DCE registry 102. The ...each foreign registry includes the foreign registry password propagation binding ERA, which allows the password synchronization server to locate the foreign registry for purposes of password propagation.

Accounts for each password strength server 110 are also provided. Again, only one password strength server 110 is illustrated in Figure 3, but accounts for each of the password strength servers X-Z may be included. Each such account contains the secondary password management binding ERA, which is used by the password synchronization server to locate the password strength server for a given client.

Lastly, the DCE registry further includes an account for the password synchronization server 106 and comprises three ERAs. These three ERA's include password propagation enable, password propagation... ..interval, and foreign registry. Each of these ERAs will be described in greater detail below.

Synchronization of passwords between the DCE registry and a foreign registry requires a solution that satisfies... ..definitions in its local registry. This function is herein subsequently referred to as the "password synchronization pull" function.

2) Provides a way for a foreign registry to receive automatically password changes... ..changed in the DCE registry. This function is herein subsequently referred to as the "password synchronization push" function.

Figure 4 illustrates the password synchronization pull function, a set of operations that service a foreign registry wishing to populate its... ..user's password and this update is propagated to the new registry via the password synchronization push function described in Figure 5. The flow diagram of Figure 4 depicts a solution to this problem in that it modifies the password synchronization server of Figure 3, which is represented in Figure 3A with respect to the elements used... ..trigger." Foreign registry attempts to read this ERA result in a query to the password synchronization server that subsequently returns the password for that user. It is noted that although DCE uses ERAs to associate password information with a user's account, any mechanism or security server that associates the user's password with the user's account can also be used. For example, the security

server may have fields or use pointers for association.

A description of the operations represented in Figure 4 follows. In block 410, a password update request is issued to DCE security server 104. In this case, client W requests account creation or password change for client W's account. In Block 412, DCE security server 104 upon the request retrieves the required routing information contained in the password management binding ERA associated with client Y's account and uses it to locate the password synchronization server 106. Next, in block 414, the security server 104 sends a clear text password and client Y identity to the password synchronization server 106. Afterwards, in block 416, password synchronization server 106 encrypts this information and stores it in its password repository 112. Afterwards, in block 418, the password synchronization server 106 returns a "complete" message to the security server 104.

The security server 104, in block 420 encrypts the password and stores it in the client W account within DCE registry 102. Afterwards, the security server returns a "success" message to client W (block 422).

At any time following the completion... ..Z 108 can request a copy of client W's password by requesting the security server to retrieve the value of the plain-text password ERA associated with client W's account as shown in block 424. In block 426, the security server retrieves the plain-text password ERA value from client W's account. Then, in block 428, the security server returns this value to foreign registry Z. In block 430, foreign registry Z requests the security server to return the value of foreign registry Z's foreign registry password propagate binding ERA. In block 432, the security server retrieves the value of the foreign registry password propagate binding ERA and then returns this... ..then uses the value returned in step 428 as routing information to locate the password synchronization server 106 as well as uses the value returned in step 434 to specify the type... ..W's password while it is "on the

wire" or being transmitted between the password synchronization server and foreign registry Z that occurs in a later step. Further, according to block 436, the foreign registry Z then requests client W's password from the password synchronization server 106. In block 438, the password synchronization server 106 requests the security server 104 to return the value(s) client W's foreign registry ERA. In block 440, the security server 104 retrieves this value(s) from DCE registry 102 and then returns the value(s) to the password synchronization server in block 442.

In block 444, the password synchronization server 104 inspects ...values. This is an access control check. If the value is not identified, the password synchronization server returns a "error" message in block 446. Otherwise, in block 448, if the value is identified, the password synchronization server retrieves client W's password from the password repository and decrypts it. Next, in block 450, the password synchronization server returns client W's password to foreign registry Z.

Figure 5 illustrates the password synchronization push function, a set of operations that automatically propagate password changes for selected users (accounts... ..description of the operations represented in Figure 5 follows and is further represented by the server in Figure 3B, which is a rendition of Figure 3, but with only those elements used in the synchronization push operation. Beginning in Block 510, the client of Y first requests account creation or a password change for client Y's account. In Block 512, the security server retrieves the routing information contained in the password management binding ERA associated with client Y's account and uses it to locate to the password synchronization server. In Block 514, the security server then sends the clear text password and client Y identity to the password synchronization server.

In Block 516, the password synchronization server then requests the DCE security server to provide ERA information including the foreign registry ERA for client Y from the client Y account as well as all ERAs for the password synchronization server and foreign registry Z accounts.

In Block 518, the security server retrieves the requested ERA information shown in DCE registry 102. In Block 520, the security server returns the ERA information to the password synchronization server, which then in Block 522 validates that client Y foreign registry ERA content is also contained in the password synchronization server's foreign registry ERA. This is a "access control" feature in that it prevents one... ..is allowed to define the contents of the foreign registry ERA associated with the password synchronization server account.

In Block 524, the password synchronization server uses the content of the client Y foreign registry ERA to request the foreign registry... ..from the foreign registry Z account.

Upon completion of locating the foreign registries for each server X-Z, the system, in Block 528 and through the password synchronization server returns a "complete" message to the security server 104. In Block 530, the security server 104 encrypts the password and then stores it in the client Y account. Afterwards, in Block 532, the security server returns a "complete" message to client Y.

In Block 534, if the password propagate enable ERA indicates that it is "enabled," the password synchronization server, in block 536, sends the password to the foreign registry/server Z. Otherwise, the system returns. Next, in block 538, the foreign registry Z returns a synchronization complete status, failure status, or fails to respond if the network is down. If the... ..then returns; otherwise, the system proceeds to block 540. In block 540, if the password synchronization server receives failure status or times out awaiting a response, it re-queues this attempt for... ..expired and then returns.

Native OSF DCE provides for the support of multiple password strength servers. When the DCE registry receives a request to define or change the password for a... ..ERA associated with the requesting user to locate what it believes is a password strength server, sends the user identity and plain text password to that server, and then awaits a yes or no decision from the server. Design of the password synchronization function must not (and does not) require any change to the DCE security server, as this server may be implemented by various vendors on any platform supporting DCE. This invention cannot compel these vendors to make changes to their implementations of the

security server and must therefore preserve the operation of this interface between the security server and what it believes to be a password strength server. Design of the password synchronization server takes advantage of this interface, substituting itself in the position of a password strength server, while not perturbing the security server's belief it is a password strength server. While this is an effective approach to acquiring the plain text password, thus enabling the password synchronization server to propagate it to foreign registries, it creates a problem when one desires, for the... ..underscore)binding ERA for a given account that is recognized by a native OSF Security server (and as just discussed, this invention cannot change security server behavior). Therefore, this one interface must support the functions of both password synchronization and password strength checking. Were it not necessary to preserve support for multiple password strength servers, each with its own customer-tailored set of composition rules, this constraint would not pose a significant problem. In this case, the password synchronization server 106 could, within its own processing steps, perform both propagation and strength checking. This not being the case, a different approach must be undertaken. Accordingly, in order to support customer-tailored password strength servers, password strength server 110 is added to interface with the password synchronization server 106.

With the addition of password strength server 110, the password synchronization server is able to route password change requests received from DCE registry 102 to the strength server 110. It then fields the password strength server's response, indicating password validity or invalidity, and returns the response to the DCE registry... ..password strength ERA that must be created to contain the name of the password strength server that performs password composition checking for a given user. Also for each password strength server record in the DCE registry, a new data item, the secondary password management binding ERA is required. It contains information that allows the password synchronization server to locate the password strength server for a given user. In order to implement the use of the password strength server while also providing password synchronization among the foreign registries, the system implements the steps illustrated in Figure 6.

Figure 6 depicts the steps that extend password synchronization to support password strength servers that may be customer-tailored according to a user's particular needs. Further, Figure 3C is a companion diagram... ..utilized during strength checking. Note that, as in the native OSF implementation of password strength server support, only one password strength server can be associated with a given user, though different users may be configured to use different password strength servers. In block 610, client W 114 requests account creation or password change for the client... ..to the first step in figure 5, block 510. Next, in block 612, the security server 104 retrieves the routing information contained in the password management binding ERA associated with the client Y's account and uses it to locate password synchronization server 106, which again is similar to that of block 512 in Figure 5. In block 614, the security

server sends a clear text password and client Y identity to the password synchronization server 106. In block 616, the password synchronization server requests the security server to return the value of the password strength ERA associated with the client Y's account and uses this value to identify the account for password strength server X110. The password synchronization server then requests in block 618 that the security server return the value of the secondary password management binding ERA associated with this account.

The security server 104 retrieves the requested ERA values and then returns these values to the password synchronization server in block 622. In block 624, the password synchronization server uses the information in the secondary password management binding as routing information to locate the password strength server X110 and requests this server perform its strength checking function.

In block 626, the password strength server X110 performs the strength checking function and then returns either a complete signal or an error signal to the password synchronization server. Then, in block 628, the password synchronization server returns the complete or error message of block 626 to the security server 104. In block 632, if an "error" message is noted, the security server returns the error message to client W 114. If the complete message

is sent, the security server encrypts the password, stores it in client W's account in block 630, and returns "complete" to the client.

PASSWORD SYNCHRONIZATION SERVER REQUIREMENTS

The password synchronization function meets the following requirements. First, synchronization causes passwords changed by DCE users to be propagated as plain-text passwords to any... ..propagation is subject to an administrator-controllable interval. This operation is referred to as "password synchronization push."

Second, the password synchronization server enables password synchronization without modifications to the DCE Registry code, such that any vendor's DCE 1.1 Security Server can be used to support the password synchronization functions.

Third, the password synchronization server preserves support for the password composition server (pwd(underscore)strength) function provided by OSF DCE 1.1, and allows password strength checking to occur independently of whether a user's password is to be synchronized.

Fourth, the password synchronization server supports plain-text password retrieval by a foreign registry upon foreign registry request. This operation is also called "password synchronization pull."

Fifth, the password synchronization server provides secure transmission of plain-text passwords parameters within the constraints imposed by client and server access to application data encryption functions.

Sixth, the password synchronization server provides fault-tolerance. If the password synchronization server goes down, this prevents passwords from being updated in DCE accounts of interest to foreign... ..thus precludes an out-of-sync condition. To cover the case in which the password synchronization server goes down before having emptied its in-memory propagation queues, disk-mirroring of the queues is employed. If a foreign registry is down while password changes are made, the password synchronization server maintains the required state information to be able to attempt propagation when/if the foreign... ..comes back on line.

Seventh, in providing Fault-tolerance (the sixth feature above), the password synchronization server recovers plain text passwords from disk storage. It protects these passwords while they are disk... ..these requirements may be substantially modified, changed, or eliminated when implemented in a non-DCE server environment. In the non-...implemented if the overall system is protected from hackers or other unwelcome guests.

The password synchronization server can be implemented in either a AIX or OS/2 platform, or any other type... ..over the wire encryption" protection of plain-text passwords between some nodes participating in password synchronization is accommodated. For example, lack of protection between two nodes can occur under either of... ..encryption algorithm called Commercial Data Masking Facility (CDMF) which can be used by the password synchronization function to provide "over the wire encryption" in the absence of access to DES data... ..wire encryption" protection for plain-text passwords under the following circumstances:

- a) either the security server or the password strength server is unable to support data privacy (occurs on normal composition check).
- b) either a client or the password strength server is unable to support data privacy (occurs when the client attempts to request a password strength server-generated password).

All customer-provided strength servers that may potentially cohabit in a cell containing a password synchronization server must supplement a check made in the OSF DCE sample code, whereby a strength check... ..caller is either "dce-rgy" or "pwsync." (These are the principals used by the security server and password synchronization servers when acting as a client to password strength servers.)

It is not necessary to preserve the ability for a strength server to know the real client identity on a "generate password" request, whenever such requests are actually relayed to it via the password synchronization server.

When a principal is configured for password synchronization, but not for password strength checking, the minimum set of registry policy checks on password... (alphanumeric, minimum length) are not enforced.

All foreign registries configured for support by the password synchronization server must reside in the same DCE cell as the latter.

All foreign registries should verify that the client identity initiating any "password push" RPC is that of the password synchronization server (principal "pwsync"), to thwart any attacker's attempt to spoof the password synchronization server.

The password synchronization server cannot be replicated.

OVERVIEW OF OSF DCE 1.1 PASSWORD STRENGTH CHECKING

A review of... provide a foundation for understanding how password strength checking architecture can be leveraged to implement password synchronization.

By itself, the DCE security server provides a limited number of password composition rules. This includes minimum length, whether the password... spaces. DCE 1.1 extends the capability of password composition checking by having the security server call out to a password strength-checking server whenever a password change request is received. This server also serves as a "password generator," and can be modified by customers to enforce any desired password composition rules.

The security server checks are not enforced if a strength server exists for a user; however, the behavior of the default (OSF-provided) strength server is to enforce these same checks. Strength servers can be written to ignore enforcement of such specific checks.

To support this new functionality... MGMT(underscore)BINDING: This ERA attaches to a normal user principal and defines the strength server to be called by `secd` whenever an attempt is made to alter this user's password. This ERA is also used by a "change password program" client and defines the server to invoke to obtain a generated password. This ERA is of the "binding" encoding type, containing this information:

- + service principal name of strength server
- + CDS namespace entry where server exports bindings --OR-- contains a string binding instead
- + RPC authentication levels to be used in communicating with the server.

An example of how this displays, using the `dcecp` convention, is:

`(pwd(underscore)mgmt(underscore... mgmt/pwd(underscore)strength))`

Password checking and password generation always take place in the same server, hence the need for only one ERA to define the location of this server.

0 PWD(underscore)VAL(underscore)TYPE: This ERA is also attached to normal user principals... itself.

+ USER(underscore)SELECT (1): This means that password strength checking is performed by the server named in `PWD(underscore)MGMT(underscore)BINDING` whenever `secd` receives a request to create or ...prompts for a new password should first display a generated password obtained from the strength server named by `PWD(underscore)MGMT(underscore)BINDING`. The user is free to type this or... Upon receipt of a request to change an account password, `secd` always invokes the strength server check. If the user types the generated password, the strength server algorithm treats the new password as if the user concocted it, and can potentially reject... prompts for a new password should first display a generated password obtained from a strength server named by `PWD(underscore)MGMT(underscore)BINDING`. The user is required to supply this password as confirmation. Subsequently, `secd` invokes the strength server check, which will succeed only if the strength server

sees, within a cache it maintains, that the password had recently been supplied to this... ..text password is passed as either an input or an output parameter to the strength server.

Rsec(underscore)pwd(underscore)mgmt(underscore)str(underscore)chk passes plain text as an input... ..output. Control of the degree of protection on the password is accomplished by the security server's use of the protection level specified as one of the parameters in the authentication... ..level can only be set to packet privacy or packet integrity.

Part of the password synchronization implementation is a modification of OSF's DCE dcecp program and the corresponding underlying api... ..to obtain a DES export license, IBM DCE support for the password strength and password synchronization functions provides a potential advantage over other vendors' implementations in that it will protect passwords on the wire in some circumstances where other implementations are unable to do so.

PASSWORD SYNCHRONIZATION OVERVIEW

The implementation of the password synchronization according to the present invention is now given. The system adds to or replaces the server that receives the callouts for password strength checking with a server that propagates passwords (a "push") to foreign registries instead. Before actually doing a "push," this server contacts the real strength server for password composition checking. In addition, this new password synchronization server fields requests to retrieve plain-text passwords (a "pull") from foreign registries for a specified... ..before when attached to a user principal, except that for a user for which password synchronization is relevant, the information within this ERA will point to 'pwsync,' and not a real strength server. Users for which strength checking is relevant, but not synchronization, have the ERA set to point directly to a strength server.

An instance of this ERA also is attached to the 'pwsync' principal itself. This serves... ..tool need not prompt for information necessary to have a new user participate in password synchronization; the data can be read from the value of this ERA as stored on the... ..underscore)VAL(underscore)TYPE: Same meanings as described above. Note that if 'pwsync' is the server pointed to by PWD(underscore)MGMT(underscore)BINDING, 'pwsync' does not re-read the PWD... ..the strength check RPC to in turn pass on such requests to the real strength server.

PWD(underscore)VAL(underscore)TYPE must be set to 1 or greater if any password synchronization is to occur for a user principal. PWD(underscore)VAL(underscore)TYPE should not exist... ..for foreign registries. It contains binding and authentication information for use by the password sync server in propagating passwords to the foreign registry on a "push."

0 FOREIGN(underscore)REGISTRY: This... ..normal user's principal object. It contains one or more character strings, identifying the DCE server principal name(s) of a foreign registry(s). It thus represents a registry(s) to... ..password should be propagated. It is also used in a "pull" operation by the password synchronization service as an access control mechanism. This is accomplished by insuring that the DCE identity...MGMT(underscore)BINDING: This ERA is associated with the principal object for a password strength server. It contains binding and authentication information for use by the password synchronization server in relaying requests for password generation and password checking. Since pwsync is the direct recipient... ..that pwsync can route such requests to an actual generator/checker. Such callouts to strength servers occur before any attempt to propagate a new password to foreign registries. 0 PASSWORD(underscore... ..a normal user's principal object. It contains a character string value, namely the DCE server principal name of a strength server. Thus, it denotes the strength server to be involved in any of this user's password changes.

As can easily be... ..underscore)STRENGTH value serves as a "key" during a password change, directing the password sync server to the proper service principal from which to query the binding information as stored in the SECONDARY(underscore)PWD(underscore)MGMT(underscore)BINDING era.

0 AVAILABLE(underscore)STRENGTH(underscore)SERVER: This multi-valued ERA is stored on the pwsync principal only. It contains the names of all installed strength servers. Tools that need to know the names of all strength servers as candidates for attaching a single-valued PASSWORD(underscore)STRENGTH era on a new principal can obtain the list of candidate servers by reading the value of AVAILABLE(underscore)STRENGTH(underscore)SERVER from the pwsync principal.

The PWSYNC facility also inspects AVAILABLE(underscore)STRENGTH(underscore)SERVER contents as attached to the "pwsync" principal, as a validity check on the legal values... ..(underscore)STRENGTH ERAs attached to user principals. Hence, the presence of AVAILABLE(underscore)STRENGTH(underscore)SERVER is mandatory, not optional.

0 PLAINTEXT(underscore)PASSWORD: This ERA is a query trigger associated... ..value of this ERA. The request will be rerouted from seed to the password sync server, by virtue of the query trigger stored in the schema definition having binding and authentication... ..is moot if PW(underscore)PROPAGATE(underscore)ENABLE is marked to disable propagation altogether.

PASSWORD SYNCHRONIZATION IMPLEMENTATION

The basic structure of the password synchronization server consists of a main routine that performs setup tasks typical of DCE servers and then 'listens' for requests, a key management thread, an identity refresh thread, a Pull... ..and the rsec(underscore)pwd(underscore)mgmt(underscore)str(underscore)chk RPC from the security server. The latter RPC, though, as implemented in pwsync is now basically a way-station. Rather... ..perform the actual core functionality itself, it now forwards the request to the password strength server in effect for the subject principal.

New functionality consists primarily of the functions that support "push" and "pull" operations requested by foreign registries wishing to maintain password synchronization with the DCE registry.

rsec(underscore)pwd(underscore)mgmt(underscore)str(underscore)chk serves as... ..propagation of a user's plain-text DCE password to interested and authorized foreign registry servers whenever the password is changed in DCE. Such "pushes" do not occur as an integral part of the rsec(underscore)pwd(underscore)mgmt(underscore)str(underscore)chk processing, as one might initially suspect. Delaying the return to seed is undesirable from a human factors perspective. Thus, the data is instead stored in an in-memory queue for later processing by separate propagation threads. (The first propagation attempt is immediate, with retries taking place at synchronization server's propagation thread performs a "push" operation, it is acting as a client of the foreign registry exporting the relevant "push" server interface (rsec(underscore)pwd(underscore)propagate).

"Pull" refers to the password synchronization server's role as trigger server with reference to the PLAINTEXT(underscore)PASSWORD extended registry attribute. Foreign registries' queries of the... ..this database must be afforded extra protection. Such protection is modeled after the DCE Security Server's master-key encryption of sensitive Registry data. (The same protection is afforded to passwords... ..is less crucial because the data there is expected to be short-lived.)

The Password Synchronization server consists of a single process with a main thread and five auxiliary threads. Flow diagrams detailing the processing steps for two of these threads and for the rsec(underscore)pwd(underscore)mgmt(underscore)str(underscore)chk RPC interface servicing calls from the security server follow.

1) Main Thread: Performs initialization functions required by all such DCE application servers, e.g., registers bindings and interfaces in the Cell Directory namespace, creates all auxiliary threads, "listens" on RPC interfaces to service calls from the security server and from foreign registry clients attempting to "pull" a password for a specific account.

- 2) Password Synchronization Server Identity Refresh Thread: Re-establishes this server's DCE identity by "logging in" as the server on a periodic basis. (standard practice and method for all DCE application servers).
- 3) Password Synchronization Server Key Refresh Thread: Changes the account password for this server on a periodic basis, which is well known in the practice and method for all DCE application servers.
- 4) Password Synchronization Server Pull RPC and Pull Data Base Pruning Thread: At initialization, the password synchronization server main thread initializes the Memory Pull Data Base from the Disk Pull Data Base and eliminates all entries for each... If found, the contents of the foreign(underscore)registry ERAs associated with both the password synchronization server's and requested user's account are inspected to determine if the requesting foreign registry... each user present that are superseded by another entry for the same user.
- 5) Password Synchronization Server RPC from Security Server: When the main thread receives a call from the security server via this RPC, the processing steps depicted in Figure 7 are performed.
- 6) Password Synchronization Server Propagation Queue Thread: When the password synchronization server sets Signal "P" (shown in Figure 7), this thread is awakened to perform the processing steps depicted in Figure 8.
- 7) Password Synchronization Server Retry Queue Thread: When the Propagation Queue thread sets Signal "R" (shown in Figure 8), this thread is awakened to perform the processing steps depicted in Figure 9.

Figure 7 depicts a flow chart of user password processing from the security server. First, in block 710, the password synchronization servers receives the user ID and password from the security server. Then, in block 712, the synchronization server determines if the password strength ERA is present for the user and if so proceeds to block 714; otherwise, the security synchronization server proceeds to block 720 described below. In block 714, the synchronization server uses the password strength ERA to identify the appropriate password strength server associated with the user and its account and then reads its secondary password management binding ERA. After completing block 714, the synchronization server proceeds to block 716 where the password strength server account and secondary management binding ERA are used to locate the password strength server and call the strength server with the user ID/password.

In block 718, the password strength server returns a message of whether its composition check was successful. If so, if the system... message as being a "success." If an error has occurred or if the password strength

server composition check failed, then, in block 732, the system sets the status message as "failure... the system then proceeds to block 742, described below.

In block 730, a password strength server determines if a propagation has been enabled as a function of the value of the... otherwise, in the system proceeds to block 742 described below. In block 734, the strength server stores the entry in memory in the ...blocks 736 and 738, respectively. The system then proceeds to block 740 where the strength server sets the signal "P" for propagation queue thread described above. Afterwards, and if the propagation... set in block 730, the system, in block 742, returns the status to the security server.

The password synchronization server propagation queue thread operates according to the steps illustrated in Figure 8. The password synchronization server, in block 810, initialized a memory propagation queue from the disk propagation queue, which are... if not returns to block 812; otherwise, the system proceeds to block 818 where the synchronization server processes the first or next entry in queue 736. Next, in block 820, the system retrieves... against the list of foreign registries found in the foreign registry ERA for the password synchronization server. Once the validation has been completed, the particular foreign registry on the propagation queue 736... the system determines whether there are additional foreign registries for this particular user ID being processed and if so returns to block 822 et seq. Otherwise, if no more foreign registries... more entries in the memory propagation queue, the system returns to block 818 and the process starts over again for the next entry in the propagation queue; otherwise,

the system proceeds to block 842, where the synchronization server deletes all entries in memory propagation queue 736 and disk propagation queue 738. Afterwards, the system then returns to block 812, awaiting the settling of signal 'P'.

All propagation retries are processed according to the method depicted in Figure 9. Figure 9 depicts a flow chart of the password synchronization server retry queue thread according to the present invention. First, in block 910, the password synchronization server initializes the memory retry queue from the disk retry queue at blocks 834 and 836, respectively. Next, the... ..920; otherwise, the system proceeds to block 938 described below.

In block 920, the system processes the first or next identified entry in queue 834 and then inspects this entry in... ..NOW COMMUNICATING." If so, it is reset; if not, it is left as is and processing proceeds to block 932.

In block 932, the synchronization server determines if any more entries are present in the memory entry queue 834 and if... ..916 for the next retry cycle.

PLAIN-TEXT PASSWORD PROTECTION

Operation of the password security synchronization (PWSYNC) 106 involves several RPCs that pass a plain-text password as a parameter. It... ..running in a USA or international environment, or whether one is operating with compatible DCE servers or some other vendor's servers. Before thinking about such scenarios, it is helpful to understand the nature of the RPCs... ..controls are used, how they are set, and whether they address all situations involving client/server communication in the password synchronization environment.

There are seven distinct points where programs will issue an RPC that should be... ..they only have to be specified once, during the installation of either pwsync, a strength server, or a foreign registry server. This is true whether the installation package is from the same or a different vendor... ..of these ERAs are dependent upon information related to foreign registries and to password strength servers. Without the aid of an administrative tool, e.g., GUI, that "remembers" or has access to the pertinent information associated with these servers when they were installed, and that has "canned" information defining what ERAs need to be... ..as FOREIGN(underscore)REGISTRY and PASSWORD(underscore)STRENGTH.

The method by which foreign registries and customer strength servers shall communicate such information to administrative tools is to attach the information onto the pwsync principal itself. These customer or vendor-written administrative tools, which use well-known OSF DCE functions to accomplish their... ..of the chart indicates whether the particular protection level is assigned to a single "client/server" pair, or whether it can be shared between multiple pairs. For example, the second and... ..there is a separate ERA instance for each and every foreign registry or password strength server communicated with by pwsync. Thus, if some foreign registries require different protection levels, these would... ..in this column indicates that a specified protection level is not unique across all client/server combinations. For example, item 4 has the problem that all "pulls" from foreign registries are... ..made to take advantage of the maximum degree of protection afforded by a particular client/server pair:

1) The sec(underscore)rgy(underscore)attr(underscore)lookup(underscore)* APIs modified to recognize...principal who wishes to change his password from this machine be provided with a password synchronization server-generated password, or, downgrade the protection level specified in PWD(underscore)MGMT(underscore)BINDING for... ..wire" for this principal.

Schema definitions for the new ERAs are required for this password synchronization implementation. The value of the "reserved" flag is set TRUE for all of these ERAs... ..to foreign registries. All foreign registries must reside in the same cell as the password synchronization server. The username and time of DCE registry update are also provided. If the foreign registry... ..cannot handle the notification at the present time, but would like for the password sync server to send the event in a later propagation cycle, then an unsuccessful status, of any... ..to those skilled in the art.

In the case of the password pull database, the server takes on the added overhead of locking macros and software, patterned after the DCE security server (secd) and adapted for the password synchronization server's use. Use of the locking paradigm is deemed necessary so that when the pull... ..threads.

As indicated above, aspects of this invention pertain to specific "method functions" implementable on computer systems. In an alternate embodiment, the invention may be implemented as a computer program product for use with a computer system. Those skilled in the art should readily appreciate that programs defining the functions of the present invention can be delivered to a computer in many forms, which include, but are not limited to: (a) information permanently stored on non-writable storage media (e.g. read only memory devices within a computer such as ROM or CD-ROM disks readable by a computer I/O attachment); (b) information alterably stored on writable storage media (e.g. floppy disks and hard drives); or (c) information conveyed to a computer through communication media, such as a network, and telephone networks, via modem. It should be understood, therefore, that such media, when carrying computer readable instructions that direct the method functions of the present invention, represent alternate embodiments of...

Specification: ...the secondary password management binding ERA is required. It contains information that allows the password synchronization server to locate the password strength server for a given user. In order to implement the use of the password strength server while also providing password synchronization among the foreign registries, the system implements the steps illustrated in Figure 6.

Figure 6 depicts the steps that extend password synchronization to support password strength servers that may be customer-tailored according to a user's particular needs. Further, Figure 3C is a companion diagram... ..utilized during strength checking. Note that, as in the native OSF implementation of password strength server support, only one password strength server can be associated with a given user, though different users may be configured to use different password strength servers. In block 610, client W 114 requests account creation or password change for the client... ..to the first step in figure 5, block 510. Next, in block 612, the security server 104 retrieves the routing information contained in the password management binding ERA associated with the client Y's account and uses it to locate password synchronization server 106, which again is similar to that of block 512 in Figure 5. In block 614, the security server sends a clear text password and client Y identity to the password synchronization server 106. In block 616, the password synchronization server requests the security server to return the value of the password strength ERA associated with the client Y's account and uses this value to identify the account for password strength server X 110. The password synchronization server then requests in block 618 that the security server return the value of the secondary password management binding ERA associated with this account.

The security server 104 retrieves the requested ERA values and then returns these values to the password synchronization server in block 622. In block 624, the password synchronization server uses the information in the secondary password management binding as routing information to locate the password strength server X 110 and requests this server perform its strength checking function.

In block 626, the password strength server X 110 performs the strength checking function and then returns either a complete signal or an error signal to the password synchronization server. Then, in block 628, the password synchronization server returns the complete or error message of block 626 to the security

server 104. In block 632, if an "error" message is noted, the security server returns the error message to client w 114. If the complete message is sent, the security server encrypts the password, stores it in client w's account in block 630, and returns "complete" to the client.

PASSWORD SYNCHRONIZATION SERVER REQUIREMENTS

The password synchronization function meets the following requirements. First, synchronization causes passwords changed by DCE users to be propagated as plain-text passwords to any... ..propagation is subject to an administrator-controllable interval. This operation is referred to as "password synchronization push."

Second, the password synchronization server enables password synchronization without modifications to the DCE Registry code, such that any vendor's DCE 1.1 Security Server can be used to support the password synchronization functions.

Third, the password synchronization server preserves support for the password composition server (pwd(underscore)strength) function provided by OSF DCE 1.1, and allows password strength checking to occur independently of whether a user's password is to be synchronized.

Fourth, the password synchronization server supports plain-text password retrieval by a foreign registry upon foreign registry request. This operation is also called "password synchronization pull."

Fifth, the password synchronization server provides secure transmission of plain-text passwords parameters within the constraints imposed by client and server access to application data encryption functions.

Sixth, the password synchronization server provides fault-tolerance. If the password synchronization server goes down, this prevents passwords from being updated in DCE accounts of interest to foreign... ..thus precludes an out-of-sync condition. To cover the case in which the password synchronization server goes down before having emptied its in-memory propagation queues, disk-mirroring of the queues is employed. If a foreign registry is down while password changes are made, the password synchronization server maintains the required state information to be able to attempt propagation when/if the foreign... ..comes back on line.

Seventh, in providing Fault-tolerance (the sixth feature above), the password synchronization server recovers plain text passwords from disk storage. It protects these passwords while they are disk... ..these requirements may be substantially modified, changed, or eliminated when implemented in a non-DCE server environment. In the non-DCE case, it is recommended that the requirements be followed, but... ..implemented if the overall system is protected from hackers or other unwelcome guests.

The password synchronization server can be implemented in either a AIX or OS/2 platform, or any other type... ..over the wire encryption" protection of plain-text passwords between some nodes participating in password synchronization is accommodated. For example, lack of protection between two nodes can occur under either of... ..encryption algorithm called Commercial Data Masking Facility (CDMF) which can be used by the password synchronization function to provide "over the wire encryption" in the absence of access to DES data... ..wire encryption" protection for plain-text passwords under the following circumstances:

- a) either the security server or the password strength server is unable to support data privacy (occurs on normal composition check).
- b) either a client or the password strength server is unable to support data privacy (occurs when the client attempts to request a password strength server-generated password).

All customer-provided strength servers that may potentially cohabit in a cell containing a password synchronization server must supplement a check made in the OSF DCE sample code, whereby a strength check... ..caller is either "dce-rgy" or "pwsync." (These are the principals used by the security server and password synchronization servers when acting as a client to password strength servers.)

It is not necessary to preserve the ability for a strength server to know the real client identity on a "generate password" request, whenever such requests are actually relayed to it via the password synchronization server.

When a principal is configured for password synchronization, but not for password strength checking, the minimum set of registry policy checks on password... ..alphanumeric, minimum length) are not enforced.

All foreign registries configured for support by the password synchronization server must reside in the ...verify that the client identity initiating any "password push" RPC is that of the password synchronization server (principal "pwsync"), to thwart any attacker's attempt to spoof the password synchronization server.

The password synchronization server cannot be replicated.

OVERVIEW OF OSF DCE 1.1 PASSWORD STRENGTH CHECKING

A review of... ..provide a foundation for understanding how password strength checking architecture can leveraged to implement password synchronization.

By itself, the DCE security server provides a limited number of password composition rules. This includes minimum length, whether the password... ..spaces. DCE 1.1 extends the capability of password composition checking by having the security server call out to a password strength-checking server whenever a password change request is received. This server also serves as a "password generator," and can be modified by customers to enforce any desired password composition rules.

The security server checks are not enforced if a strength server exists for a user; however, the behavior of the default (OSF-provided) strength server is to enforce these same checks. Strength servers can be written to ignore enforcement of such specific checks.

To support this new functionality... ..MGMT(underscore)BINDING: This ERA attaches to a normal user principal and defines the strength server to be called by secd whenever an attempt is made to alter this user's password. This ERA is also used by a "change password program" client and defines the server to invoke to obtain a generated password. This ERA is of the "binding" encoding type, containing this information:

- + service principal name of strength server
- + CDS namespace entry where server exports bindings --OR-- contains a string binding instead
- + RPC authentication levels to be used in communicating with the server.

An example of how this displays, using the dcecp convention, is:

Password checking and password generation always take place in the same server, hence the need for only one ERA to define the location of this server.

0 PWD(underscore)VAL(underscore)TYPE: This ERA is also attached to normal user principals... ..itself.

+ USER(underscore)SELECT (1): This means that password strength checking is performed by the server named in PWD(underscore)MGMT(underscore)BINDING whenever secd receives a request to create or... ..prompts for a new password should first display a generated password obtained from the strength server named by PWD(underscore)MGMT(underscore)BINDING. The user is free to type this or... ..Upon receipt of a request to change an account password, secd always invokes the strength server check. If the user types the generated password, the strength server algorithm treats the new password as if the user concocted it, and can potentially reject... ..prompts for a new password should first display a generated password obtained from a strength server named by PWD(underscore)MGMT(underscore)BINDING. The user is required to supply this password as confirmation. Subsequently, secd invokes the strength server check, which will succeed only if the strength server sees, within a cache it maintains, that the password had recently been supplied to this... ..text password is passed as either an input or an output parameter to the strength server.

Rsec(underscore)pwd(underscore)mgmt(underscore)str(underscore)chk passes plain text as an input... ..output. Control of the degree of protection on the password is accomplished by the security server's use of the protection level specified as one of the parameters in the authentication... ..level can only be set to packet privacy or packet integrity.

Part of the password synchronization implementation is a modification of OSF's DCE dcecp program and the corresponding underlying api... ..to obtain a DES export license, IBM DCE support for the password strength and password synchronization functions provides a potential advantage over other vendors' implementations in that it will protect passwords on the wire in some circumstances where other implementations are unable to do so.

PASSWORD SYNCHRONIZATION OVERVIEW

The implementation of the password synchronization according to the present invention is now given. The system adds to or replaces the server that receives the callouts for password strength checking with a server that propagates passwords (a "push") to foreign registries instead. Before actually doing a "push," this server contacts the real strength server for password composition checking. In addition, this new password synchronization server fields requests to retrieve plain-text passwords (a "pull") from foreign registries for a specified...before when attached to a user principal, except that for a user for which password synchronization is relevant, the information within this ERA will point to 'pwysnc,' and not a real strength server. Users for which strength checking is relevant, but not synchronization, have the ERA set to point directly to a strength server.

An instance of this ERA also is attached to the 'pwysnc' principal itself. This serves... ..tool need not prompt for information necessary to have a new user participate in password synchronization; the data can be read from the value of this ERA as stored on the... ..underscore)VAL(underscore)TYPE: Same meanings as described above. Note that if 'pwysnc' is the

server pointed to by PWD(underscore)MGMT(underscore)BINDING, 'pwysnc' does not re-read the PWD... ..the strength check RPC to in turn pass on such requests to the real strength server.

PWD(underscore)VAL(underscore)TYPE must be set to 1 or greater if any password synchronization is to occur for a user principal. PWD(underscore)VAL(underscore)TYPE should not exist... ..for foreign registries. It contains binding and authentication information for use by the password sync server in propagating passwords to the foreign registry on a "push."

0 FOREIGN(underscore)REGISTRY: This... ..normal user's principal object. It contains one or more character strings, identifying the DCE server principal name(s) of a foreign registry(s). It thus represents a registry(s) to... ..password should be propagated. It is also used in a "pull" operation by the password synchronization service as an access control mechanism. This is accomplished by insuring that the DCE identity...

...MGMT(underscore)BINDING: This ERA is associated with the principal object for a password strength server. It contains binding and authentication information for use by the password synchronization server in relaying requests for password generation and password checking. Since pwysnc is the direct recipient... ..that pwysnc can route such requests to an actual generator/checker. Such callouts to strength servers occur before any attempt to propagate a new password to foreign registries.

0 PASSWORD(underscore)... ..a normal user's principal object. It contains a character string value, namely the DCE server principal name of a strength server. Thus, it denotes the strength server to be involved in any of this user's password changes.

As can easily be... ..underscore)STRENGTH value serves as a "key" during a password change, directing the password sync server to the proper service principal from which to query the binding information as stored in the SECONDARY(underscore)PWD(underscore)MGMT(underscore)BINDING era.

0 AVAILABLE(underscore)STRENGTH(underscore)SERVER: This multi-valued ERA is stored on the pwysnc principal only. It contains the names of all installed strength servers. Tools that need to know the names of all strength servers as candidates for attaching a single-valued PASSWORD(underscore)STRENGTH era on a new

principal can obtain the list of candidate servers by reading the value of AVAILABLE(underscore)STRENGTH(underscore)SERVER from the pwsync principal.

The PWSYNC facility also inspects AVAILABLE(underscore)STRENGTH(underscore)SERVER contents as attached to the "pwsync" principal, as a validity check on the legal values... ..STRENGTH ERAS attached to user principals. Hence, the presence of AVAILABLE(underscore)STRENGTH(underscore)SERVER is mandatory, not optional.

0 PLAINTEXT(underscore)PASSWORD: This ERA is a query trigger associated... ..value of this ERA. The request will be rerouted from secd to the password sync server, by virtue of the query trigger stored in the schema definition having binding and authentication...is moot if PW(underscore)PROPAGATE(underscore)ENABLE is marked to disable propagation altogether.

PASSWORD SYNCHRONIZATION IMPLEMENTATION

The basic structure of the password synchronization server consists of a main routine that performs setup tasks typical of DCE servers and then 'listens' for requests., a key management thread; an identity refresh thread, a Pull... ..and the rsec(underscore)pwd(underscore)mgmt(underscore)str(underscore)chk RPC from the security server. The latter RPC, though, as implemented in pwsync is now basically a way-station. Rather... ..perform the actual core functionality itself, it now forwards the request to the password strength server in effect for the subject principal.

New functionality consists primarily of the functions that support "push" and "pull" operations requested by foreign registries wishing to maintain password synchronization with the DCE registry.

rsec(underscore)pwd(underscore)mgmt(underscore)str(underscore)chk serves as... ..propagation of a user's plain-text DCE password to interested and authorized foreign registry servers whenever the password is changed in DCE. Such "pushes" do not occur as an integral part of the rsec(underscore)pwd(underscore)mgmt(underscore)str(underscore)chk processing, as one might initially suspect. Delaying the return to secd is undesirable from a human factors perspective. Thus, the data is instead stored in an in-memory queue for later processing by separate propagation threads. (The first propagation attempt is immediate, with retries taking place at... ..out-of-sync conditions between the DCE and foreign registries. Note that when the password synchronization server's propagation thread performs a "push" operation, it is acting as a client of the foreign registry exporting the relevant "push" server interface (rsec(underscore)pwd(underscore)propagate).

"Pull" refers to the password synchronization server's role as trigger server with reference to the PLAINTEXT(underscore)PASSWORD extended registry attribute. Foreign registries' queries of the... ..this database must be afforded extra protection. Such protection is modeled after the DCE Security Server's master-key encryption of sensitive Registry data. (The same protection is afforded to passwords... ..is less crucial because the data there is expected to be short-lived.)

The Password Synchronization server consists of a single process with a main thread and five auxiliary threads. Flow diagrams detailing the processing steps for two of these threads and for the rsec(underscore)pwd(underscore)mgmt(underscore)str(underscore)chk RPC interface servicing calls from the security server follow.

1) Main Thread: Performs initialization functions required by all such DCE application servers, e.g., registers bindings and interfaces in the Cell Directory namespace, creates all auxiliary threads, "listens" on RPC interfaces to service calls from the security server and from foreign registry clients attempting to "pull" a password for a specific account.

2) Password Synchronization Server Identity Refresh Thread: Reestablishes this server's DCE identity by "logging in" as the server on a periodic basis. (standard practice and method for all DCE application servers).

3) Password Synchronization Server Key Refresh Thread: Changes the account password for this server on a periodic basis, which is well known in the practice and method for all DCE application servers.

4) Password Synchronization Server Pull RPC and Pull Data Base Pruning Thread: At initialization, the password synchronization server main thread initializes the Memory Pull Data Base from the Disk Pull Data Base and eliminates all entries for each... ..If found, the contents of the foreign(underscore)registry ERAS associated with both the password synchronization server's and requested user's account are inspected to determine if the requesting foreign registry... ..each user present that are superseded by another entry for the same user.

5) Password Synchronization Server RPC from Security Server: when the main thread receives a call from the security server via this RPC, the processing steps depicted in Figure 7 are performed.

6) Password Synchronization Server Propagation Queue Thread: when the password synchronization server sets Signal "P" (shown in Figure 7), this thread is awakened to perform the processing steps depicted in Figure 8.

7) Password Synchronization Server Retry Queue Thread: when the Propagation Queue thread sets Signal "R" (shown in Figure 8), this thread is awakened to perform the processing steps depicted in Figure 9.

Figure 7 depicts a flow chart of user password processing from the security server. First, in block 710, the password synchronization servers receives the user ID and password from the security server. Then, in block 712, the synchronization server determines if the password strength ERA is present for the user and if so proceeds to block 714; otherwise, the security synchronization server proceeds to block 720 described below. In block 714, the synchronization server uses the password strength ERA to identify the appropriate password strength server associated with the user and its account and then reads its secondary password management binding ERA. After completing block 714, the synchronization server proceeds to block 716 where the password strength server account and secondary management binding ERA are used to locate the password strength server and call the strength server with the user ID/password.

In block 718, the password strength server returns a message of whether its composition check was successful. If so, if the system... ..message as being a "success." If an error has occurred or if the password strength server composition check failed, then, in block 732, the system sets the status message as "failure... ..the system then proceeds to block 742, described below.

In block 730, a password strength server determines if a propagation has been enabled as a function of the value of the... ..otherwise, in the system proceeds to block 742 described below. In block 734, the strength server stores the entry in memory in the propagation queue and also in the disk propagation... ..blocks 736 and 738, respectively. The system then proceeds to block 740 where the strength server sets the signal "P" for propagation queue thread described above. Afterwards, and if the propagation... ..set in block 730, the system, in block 742, returns the status to the security server.

The password synchronization server propagation queue thread operates according to the steps illustrated in Figure 8. The password synchronization server, in block 810, initialized a memory propagation queue from the disk propagation queue, which are... ..if not returns to block 812; otherwise, the system proceeds to block 818 where the synchronization server processes the first or next entry in queue 736. Next, in block 820, the system retrieves... ..against the list of foreign registries found in the foreign registry ERA for the password

synchronization server. Once the validation has been completed, the particular foreign registry on the propagation queue 736... ..the system determines whether there are additional foreign registries for this particular user ID being processed and if so returns to block 822 et seq. Otherwise, if no more foreign registries... ..more entries in the memory propagation queue, the system returns to block 818 and the process starts over again for the next entry in the propagation queue; otherwise, the system proceeds to block 842, where the synchronization server deletes all

entries in memory propagation queue 736 and disk propagation queue 738. Afterwards, the system then returns to block 812, awaiting the settling of signal 'P'.

All propagation retries are processed according to the method depicted in Figure 9. Figure 9 depicts a flow chart of the password synchronization server retry queue thread according to the present invention. First, in block 910, the password synchronization server initializes the memory retry queue from the disk retry queue at blocks 834 and 836, respectively. Next, the... ..920; otherwise, the system proceeds to block 938 described below.

In block 920, the system processes the first or next identified entry in queue 834 and then inspects this entry in... ..NOW COMMUNICATING." If so, it is reset; if not, it is left as is and processing proceeds to block 932.

In block 932, the synchronization server determines if any more entries are present in the memory entry queue 834 and if... ..916 for the next retry cycle.

PLAIN-TEXT PASSWORD PROTECTION

Operation of the password security synchronization (PWSYNC) 106 involves several RPCs that pass a plain-text password as a parameter. It... ..running in a USA or international environment, or whether one is operating with compatible DCE servers or some other vendor's servers. Before thinking about such scenarios, it is helpful to understand the nature of the RPCs... ..controls are used, how they are set, and whether they address all situations involving client/server communication in the password synchronization environment.

There are seven distinct points where programs will issue an RPC that should be... ..they only have to be specified once, during the installation of either pwsync, a strength server, or a foreign registry server. This is true whether the installation package is from the same or a different vendor... ..of these ERAS are dependent upon information related to foreign registries and to password strength servers. Without the aid of an administrative tool, e.g., GUI, that "remembers" or has access to the pertinent information associated with these servers when they were installed, and that has "canned" information defining what ERAS need to be customer strength servers shall communicate such information to administrative tools is to attach the information onto the pwsync principal itself. These customer or vendor-written administrative tools, which use well-known OSF DCE functions to accomplish their... ..of the chart indicates whether the particular protection level is assigned to a single "client/server" pair, or whether it can be shared between multiple pairs. For example, the second and... ..there is a separate ERA instance for each and every foreign registry or password strength server communicated with by pwsync. Thus, if some foreign registries require different protection levels, these would... ..in this column indicates that a specified protection level is not unique across all client/server combinations. For example, item 4 has the problem that all "pulls" from foreign registries are... ..made to take advantage of the maximum degree of protection afforded by a particular client/server pair:

1) The sec(underscore)rgy(underscore)attr(underscore)lookup(underscore)* APIS modified to recognize... ..principal who wishes to change his password from this machine be provided with a password synchronization server-generated password, or, downgrade the protection level specified in PWD(underscore)MGMT(underscore)BINDING for... ..wire" for this principal.

Schema definitions for the new ERAS are required for this password synchronization implementation. The value of the "reserved" flag is set TRUE for all of these ERAS... ..to foreign registries. All foreign registries must reside in the same cell as the password synchronization server. The username and time of DCE registry update are also provided. If the foreign registry... ..cannot handle the notification at the present time, but would like for the password sync server to send the event in a later propagation cycle, then an unsuccessful status, of any...to those skilled in the art.

In the case of the password pull database, the server takes on the added overhead of locking macros and software, patterned after the DCE security server (seed) and adapted for the password synchronization server's use. Use of the locking paradigm is deemed necessary so that when the pull... ..threads.

As indicated above, aspects of this invention pertain to specific "method functions" implementable on computer systems. In an alternate embodiment, the invention may be implemented as a computer program product for use with a computer system. Those skilled in the art should readily appreciate that programs defining the functions of the present invention can be delivered to a computer in many forms, which include, but are not limited to: (a) information permanently stored on non-writable storage media (e.g. read only memory devices within a computer such as ROM or CD-ROM disks readable by a computer I/O attachment); (b) information alterably stored on writable storage media (e.g. floppy disks and hard drives); or (c) information conveyed to a computer through communication media, such as a network, and telephone networks, via modem. It should be understood, therefore, that such media, when carrying computer readable instructions that direct the method functions of the present invention, represent alternate embodiments of...

Claims:

1. A network system server for providing password composition checking for a plurality of clients, the server comprising:
 - a main data store;
 - a security server, coupled to said main data store and said plurality of clients;
 - a password synchronization server, coupled to said security server;
 - a plurality of password strength servers, coupled to said password synchronization server, that provides password integrity among said plurality of clients so that each client maintains a password whose composition is consistent with said network server system.
2. A server according to claim 1 wherein said main data store further includes an information account for said plurality of password strength servers that binds each of said plurality of password strength servers with each of said plurality of clients.
3. A server according to claim 1 wherein each of said password strength servers is uniquely programmable with respect to performing password composition checking.
4. A network system server that provides password synchronization between a main data store and a plurality of secondary data stores, comprising:
 - a security server, coupled to said main data store;
 - a plurality of clients, coupled to said security server for accessing said main data store wherein each client maintains a unique, modifiable password;
 - a password synchronization server, coupled to said security server and to said plurality of secondary data stores; and
 - a password repository, coupled to said password synchronization server, that stores said passwords whereby one of said secondary data stores can retrieve said passwords via said password synchronization server so that each client is able to maintain a single, unique password among said plurality of secondary data stores.
5. A network system server for providing password synchronization between a main data store and a plurality of secondary data stores, the server comprising:
 - a security server, coupled to said main data store;
 - a plurality of clients, coupled to said security server for accessing said main data store wherein each client maintains a unique, modifiable password; and

a password synchronization server, coupled to said security server and to said plurality of secondary data stores, that provides password propagation synchronization to each of said secondary data stores from a user associated with one of said... ..to maintain a single, unique password among said plurality of secondary data stores.

6. A server according to any of claims 1, 4 or 5 wherein said main data store further... ..for each of said plurality of clients that includes a user's password.

7. A server according to claim 4 or claim 5 wherein said main data store further includes an information account for said password synchronization server that binds each of said plurality of secondary data stores with each of said plurality of clients.

8. A server according to claim 4 or claim 5 wherein said main data store further includes an information account for each of said secondary data stores.

9. A server according to claim 6 and claim 1 wherein said security server provides a clear-text password to said password synchronization server for subsequent routing to said password strength servers.

10. A server according to claim 6 and either claim 4 or claim 5 wherein said security server provides a clear-text password to said password synchronization server for propagating to each of said plurality of secondary data stores.

11. A server according to claim 9 or claim 10 wherein said security server includes means for encrypting said clear-text password for storage in said information account.

12. A server according to claim 4 wherein said main data store is a main registry within a... ..computing environment and each of said plurality of secondary data stores is a foreign registry server coupled to a memory store.

13. A server according to claim 4 or claim 5 further comprising:

a temporary memory and local data store, coupled to said password synchronization server and containing information supporting a propagation retry, that allows said password synchronization server to perform a propagation retry in the event of a temporary foreign registry or password synchronization server outage.

14. A server according to claim 5 wherein said password propagation is imposed on said plurality of local data stores regardless of the current password status of said local data stores.

15. A server according to claim 5 comprising:

a password repository, coupled to said password synchronization server, that stores said passwords whereby one of said local data stores can retrieve said passwords via said password synchronization server so that each client is able to maintain a single, unique password among said plurality of local data stores.

16. A server according to claim 5 comprising:

a plurality of password servers, coupled to said password synchronization server, that check password integrity among said plurality of clients so that each client maintains a... ..clients so that each client maintains a password whose composition is consistent with a network server system in which said plurality of clients operate.

19. A method for providing password synchronization between a main data store and a plurality of secondary data stores, comprising:

storing in... ..password selected by a user associated with one of a plurality of clients;

propagating password synchronization to each of said secondary data stores from said main data store so that said...
...unique password among said plurality of secondary data stores.

20. A method for providing password synchronization between a main data store and a plurality of secondary data stores, comprising:

maintaining a...

6/K/17 (Item 17 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Country	Number	Kind	Date
Type	Pub. Date	Kind	Text
...Compaq Computer Corporation (687797) 20555 SH 249 Houston, Texas 77070-2698 US	19		

Application...	19		
Available Text	Language	Update	Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: ...and 08/485,055 filed concurrently herewith.

The present invention is directed generally to data processing systems, and more particularly to a multiple processing system and a reliable system area network that provides connectivity for interprocessor and input/output... ..and communications systems to general purpose high availability commercial systems. The evolution of fault tolerant computers has been well documented (see D. P. Siewiorek, R. S. Swarz, "The Theory and Practice... ..and the Jet Propulsion laboratory began to apply fault tolerance to the development of guidance computers for aerospace applications. The 1960's also saw the development of the first AT&T electronic switching systems.

The first commercial fault tolerant machines were introduced by Tandem Computers in the 1970's for use in on-line transaction processing applications (J. Bartlett, "A NonStop Kernal," in proc. Eighth Symposium on operating System Principles, pp... ..systems were introduced in the 1980's (O. Serlin, "Fault- Tolerant Systems in Commercial Applications," Computer, pp. 19-30, August 1984). Current commercial fault tolerant systems include distributed memory multi-processors, shared-memory transaction based systems, "pair-and-spare" hardware fault tolerant systems (see R. Freiburghouse, "Making Processing Fail-safe," Mini-micro Systems, pp. 255-264, May 1982; U.S. Patent No. 4... ..system.), and triple-modular-redundant systems such as the "Integrity" computing system manufactured by Tandem Computers Incorporated of Cupertino, California, assignee of this application and the invention disclosed herein.

Most applications of commercial fault tolerant computers fall into the category of on-line transaction processing. Financial institutions require high availability for electronic funds transfer, control of automatic teller machines, and... ..telecommunications systems.

Vendors of fault tolerant machines attempt to achieve both increased system availability, continuous processing, and correctness of data even in the presence of faults. Depending upon the particular system architecture, application software ("processes") running on the system either continue to run despite failures, or the processes are automatically restarted from a recent checkpoint when a fault is encountered. Some fault tolerant systems are provided with sufficient component redundancy to be able reconfigure around failed components, but processes running in the failed modules are lost. Vendors of commercial fault tolerant systems have extended fault tolerance beyond the processors and disks. To make large improvements in reliability, all sources of failure must be addressed... ..power supplies, fans and inter-module connections.

The "NonStop," and "Integrity" architectures manufactured by Tandem Computers Incorporated, (both respectively illustrated broadly in U.S. Patent No. 4,228,496 and U... ..assigned to the assignee of this application; NonStop and Integrity are registered trademarks of Tandem Computers Incorporated) represent two current approaches to commercial fault tolerant computing. The NonStop system, as generally... ..above-identified U.S. Patent No. 4,278,496, employs an architecture that uses multiple processor systems designed to continue operation despite the failure of any single hardware component. In normal operation, each processor system uses its major components independently and concurrently, rather than as "hot backups". The NonStop system architecture may consist of up to 16 processor systems interconnected by a bus for interprocessor communication. Each processor system has its own memory which contains a copy of a message-based operating system. Each processor system controls one or more input/output (I/O) busses. Dual-porting of I/O controllers and devices provides multiple paths to each device. External storage (to the processor system), such as disk storage, may be mirrored to maintain redundant permanent data storage.

This... ..hardware, while fault recovery is the responsibility of the software.

Also, in the Nonstop multi-processor architecture, application software ("process") may run on the system under the operating system as "process-pairs," including a primary process and a backup process. The primary process runs on one of the multiple processors while the backup process runs on a different processor. The backup process is usually dormant, but periodically updates its state in response to checkpoint messages from the primary process. The content of a checkpoint message can take the form of complete state update, or... ..checkpoints were manually inserted in application programs, but currently most application code runs under transaction processing software which provides recovery through a combination of checkpoints and transaction two-phase commit protocols.

Interprocessor message traffic in the Tandem Nonstop architecture includes each processor periodically broadcasting an "I'm Alive" message for receipt by all the processors of the system, including itself, informing the other processors that the broadcasting processor is still functioning. When a processor fails, that failure will be announced and identified by the absence of the failed processor's periodic "I'm Alive" message. In response, the operating system will direct the appropriate backup processes to begin primary execution from the last checkpoint. New backup processes may be started in another processor, or the process may be run with no backup until the hardware has been repaired. U.S. Patent... ..example of this technique.

Each I/O controller is managed by one of the two processors to which it is attached. Management of the controller is periodically switched between the processors. If the managing processor fails, ownership of the controller is automatically switched to the other processor. If the controller fails, access to the data is maintained through another controller.

In addition to providing hardware fault tolerance, the processor pairs of the above-described architecture provide some measure of software fault tolerance. When a processor fails due to a software error, the backup processor frequently is able to successfully continue processing without encountering the same error. The software

environment in the backup processor typically has different queue lengths, table sizes, and process mixes. Since most of the software bugs escaping the software quality assurance tests involve infrequent data dependent boundary conditions, the backup processes often succeed.

In contrast to the above-described architecture, the Integrity system illustrates another approach... fault recovery is the logical choice since few modifications to the software are required. The processors and local memories are configured using triple-modular-redundancy (TMR). All processors run the same code stream, but clocking of each module is independent to provide tolerance... three streams is asynchronous, and may drift several clock periods apart. The streams are re-synchronized periodically and during access of global memory. Voters on the TMR Controller boards detect and mask failures in a processor module. Memory is partitioned between the local memory on the triplicated processor boards and the global memory on the duplicated TMRC boards. The duplicated portions of the... techniques to detect failures. Each global memory is dual ported and is interfaced to the processors as well to the I/O Processors (IOPs). Standard VME peripheral controllers are interfaced to a pair of busses through a Bus... the BIMs to switch control of all controllers to the remaining IOP. Mirrored disk storage units may be attached to two different VME controllers.

In the Integrity system all hardware failures... reintegrated on-line.

The preceding examples illustrate present approaches to incorporating fault tolerance into data processing systems. Approaches involving software recovery require less redundant hardware, and offer the potential for some... have been developed on other systems.

Thus, the systems described above provide fault tolerant data processing either by hardware (e.g., fail-functional, employing redundancy) or by software techniques (fail-fast... hardware). However, none of the systems described are believed capable of providing fault tolerant data processing, using both hardware (fail-functional) and software (fail-fast) approaches, by a single data processing system.

Computing systems, such as those described above, are often used for electronic commerce: electronic data interchange (EDI) and global messaging. Today's demands upon such electronic commerce, however, is demanding more and more throughput capacity as the number of users increases and... networks such as local area networks (LANS), and the like.

A key requirement for a server architecture is the ability to move massive quantities of data. The server should have high bandwidth that is scalable, so that added throughput capacity can be added... response time, latency affects service levels and employee productivity.

The present invention provides a multiple-processor system that combines both of the two above-described approaches to fault tolerant architecture, hardware redundancy and software recovery techniques, in a single system.

Broadly, the present invention includes a processing system composed of multiple sub-processing systems. Each sub-processing system has, as the main processing element, a central processing unit (CPU) that in turn comprises a pair of processors operating in lock-step, synchronized fashion to execute each instruction of an instruction stream at the same time. Each of the sub-processing systems further include an input/output (I/O) system area network system that provides redundant communication paths between various components of the larger

processing system, including a CPU and assorted peripheral devices (e.g., mass storage units, printers, and the like) of a sub-processing system, as well as between the sub-processors that may make up the larger overall processing system. Communication between any component of the processing system (e.g., a CPU and another CPU, or a CPU and any peripheral device, regardless of which sub-processing system it may belong to) is implemented by forming and transmitting packetized messages that are... responsible for choosing the proper or available communication paths from a transmitting component of the processing system to a destination component based

upon information contained in the message packet. Thus, the... peripherals, but permits it to also be used for interprocessor communications.

As indicated above, the processing system of the present invention is structured to provide fault-tolerant operation through both "fail... at a variety of points in the various data paths between the (lock-step operated) processor elements of the CPU and its associated memory. In particular, the processing system of the present invention conducts error-checking at an interface, and in a manner... little impact on performance. Prior art systems typically implement error-checking by running pairs of processors, and checking (comparing) the data and instruction flow between the processors and a cache memory. This technique of error-checking tended to add delay to the... error-checking precluded use of off-the-shelf parts that may be available (i.e., processor/cache memory combinations on a single semiconductor chip or module). The present invention performs error-checking of the processors at points that operate at slower rates, such as the main memory and I/O interfaces which operate at slower speeds than the processor-cache interface. In addition, the error-checking is performed at locations that allow detection of errors that may occur in the processors, their cache memory, and the I/O and memory interfaces. This allows simpler designs for... other data integrity checks.

Error-checking of the communication flow between the components of the processing system is achieved by adding a cyclic-redundancy-check (CRC) to the message packets that... "Good" (TPG) or "This Packet Bad" (TPB) - is appended to every packet. A maintenance diagnostic processor can use this information to isolate a link or router element that introduces an error... of topologies, so that alternate paths can be provided between any two elements of a processing system (e.g., between a CPU and an I/O device), for communication in the... so (e.g., by creating a "deadlock" condition, discussed further below).

The CPUs of a processing system are capable of operating in one of two basic modes: a "simplex mode" in... independently of the other, or a "duplex "mode in which pairs of CPUs operate in synchronized, lock-step fashion. Simplex mode operation provides the capability of recovering from faults that are... U.S. Pat. No. 4,228,496 which teaches a multiprocessing system in which each processor has the capability of checking on the operability of its sibling processors, and of taking over the processing of a processor found or believed to have failed). When operating in duplex mode, the paired CPUs both... fault tolerant platform for less robust operating systems (e.g., the UNIX operating system). The processing system of the present invention, with the paired, lock-step CPUs, is structured so that... masked (i.e., operating despite the existence of a fault), primarily through hardware.

When the processing system is operating in duplex mode, each CPU pair uses the I/O system to access any peripheral of the processing system, regardless of which (of the two, or more) sub-processor system the peripheral may be ostensibly a member of. Also, in duplex mode, message packets... message for the CPU pair (from either a peripheral device such as a mass storage unit or from a processing unit), will replicate the message and deliver it to both CPUs of the pair using synchronization methods that ensure that the CPUs remain synchronized. In effect, the duplex CPU pair, as viewed from the I/O system and other... as a single CPU. Thus, the I/O system, which includes elements from all sub-processing systems, is made to be seen by the duplex CPU pair as one homogeneous system... a multiprocessor system in which the CPU of any one is actually a pair of synchronized, lock-step CPUs.

Yet another important aspect of the present invention is that interrupts issuing... interrupts via the message packet system ensures that they will arrive at duplexed CPUs in synchronized fashion, in the same manner as I/O message packets. Interrupt message packets will contain... the system. In addition, using the same messaging system to communicate data between I/O units and the CPUs and to communicate interrupts to the CPUs preserves the ordering of I... the implementation of a technique of validating access to the memory of any CPU. The processing system, as structured according to the present invention, permits the memory of any CPU to... to handle input/output information transfers between a CPU and any other component of the processor system. Thereby, the individual processor units of the CPU are removed from the more mundane tasks of getting information from memory and out onto the TNet network, or accepting information from the network. The processor unit of the CPU merely sets up data structures in memory containing the data to be... is required, where in memory the response is to be placed

when received. When the processor unit completes the task of creating the data structure, the block transfer engine is notified to... ..response is received, it is routed to the expected memory location identified, and notifies the processor unit that the response was received.

Further aspects and features of the present invention will become... ..invention, which should be taken in conjunction with the accompanying drawings.

Fig. 1A illustrates a processing system constructed in accordance with the teachings of the present invention, and Figs. 1B and 1C illustrate two alternate configurations of the processing system of Fig. 1A, employing clusters or arrangements of the processing system of Fig. 1A;

Fig. 2 illustrates, in simplified block diagram form, the central processing unit (CPU) that forms a part of each sub-processor system of Figs. 1A - 1C;

Figs. 3A - 3D and 4A - 4C illustrate the construction of... ..the area network I/O system shown in Fig. 2;

Fig. 5 illustrates the interface unit that forms a part of the CPUs of Fig. 2 to interface the processor and memory with the I/O area network system;

Fig. 6 is a block diagram, illustrating a portion of packet receiver of the interface unit of Fig. 5;

Fig. 7A diagrammatically illustrates the clock synchronization FIFO (CS FIFO) used by the packet receiver section packet receiver shown in Fig. 6;

Fig. 7B is an block diagram of a construction of the clock synchronization FIFO structure shown in Fig. 7A;

Fig. 8 illustrates the cross-connections for error-checking outbound transmissions from the two interface units of a CPU;

Fig. 9 illustrates an encoded (8B to 9B) data/command symbol;

Fig. 10 illustrates the method and structure used by the interface unit of Fig. 5 to cross-check for errors data being transferred to the memory controllers... ..of a CPU of Fig. 2 to other (external to the CPU) components of the processing system;

Fig. 12 is a block diagram that diagrammatically illustrates the formation of an address... ..14A illustrates the logic for posting interrupt requests to queues in memory and to the processor units of the CPU of Fig. 2;

Fig. 14B illustrates the process used to form a memory address for a queue entry;

Fig. 15 is a block... ..data output constructs formed in the memory of the CPU of Fig. 2 by a processor unit, and containing data to be sent via the area I/O networks shown in Figs. 1A - 1C, and also illustrating the block transfer engine (BTE) unit of the interface unit of Fig. 5 that operates to access the data output constructs for transmission to the... ..pair of memory controllers between memory of a CPU of Fig. 2 and its interface unit for accessing from memory 72 bits of data, including two simultaneously-accessed 32-bit words... ..other for error-checking;

Fig. 19A is a simplified block diagram illustration of the router unit used in the area input/output networks of the processing systems shown in Figs. 1A - 1C;

Fig. 19B illustrates comparison on two port inputs of the router unit of Fig. 19A;

Fig. 20A is a block diagram the construction of one of the six input ports of the router unit shown in Fig. 19A;

Fig. 20B is a block diagram of the synchronization logic used to validate command/data symbols received at an input port of the router unit of Fig. 19A;

Fig. 21A is a block diagram illustration of the target port selection... ..is a block diagram illustration of one of the six output ports of the router unit shown in Fig. 19A;

Fig. 23 is an illustration of the method used to transmit identical information to a duplexed pair CPUs of Fig. 2 in synchronized fashion when the processing system is operating in lock-step (duplex) mode, using a pair the FIFOs of Fig... ..is a simplified block diagram illustrating the clock generation system of each of the sub-processing systems of Figs. 1A - 1C for developing the plurality of clock signals used to operate the various elements of that sub-processing system;

Fig. 25 illustrates the topology used to interconnect the clock generation systems of paired sub-processing systems for synchronizing the various clock signals of the pair of sub-processing systems to one another;

Fig. 26A and 26B illustrates a FIFO constant rate clock control logic used to control the clock synchronization FIFO of Figs. 8 or 20 in the situation when the two clocks used to... ..structure of the on-line access port (OLAP) used to provide access to the maintenance

processor (MP) to the various elements of the system of Fig. 1A (or those of Figs... ..the soft-flag logic used to handle asymmetric variables between the CPUs of paired sub-processing systems operating in duplex mode;

Fig. 31A shows a flow diagram, and Fig. 31B illustrates a portion of SYNC CLK, both of which are used to reset and synchronize the clock synchronization FIFOs of the CPUs and routers of the processing system of Fig. 1A that receive information from each other;

Fig. 32 is a flow... ..33A - 33D generally illustrate the procedure used to bring an one of the CPUs of processing system shown in Fig. 1A into lock-step, duplex mode operation with the other of the CPUs without measurably halting operation of the processing system; and

Fig. 34 illustrates a reduced cost architecture incorporating teachings of the invention; and... ..to the figures and, for the moment, principally Fig. 1A, there is illustrated a data processing system, designated with the reference 10, constructed according to the various teachings of the present invention. As Fig. 1A shows, the data processing system 10 comprises two sub-processor systems 10A and 10B each of which are substantially the same in structure and function... ..should be appreciated that, unless noted otherwise, a description of any one of the sub-processor systems 10 will apply equally to any other sub-processor system 10.

Continuing with Fig. 1A therefore, each of the sub-processor systems 10A, 10B is illustrated as including a central processing unit (CPU) 12, a router 14, and a plurality of input/output (I/O) packet interfaces... ..one of the I/O packet interfaces 16 will also have coupled thereto a maintenance processor (MP) 18.

The MP 18 of each sub-processor system 10A, 10B connects to each of the elements of that sub-processor system via an IEEE 1149.1 test bus 17 (shown in phantom in Fig. 1A... ..accompanying clock signal. As Fig. 1A further illustrates, TNet Links L also interconnect the sub-processor systems 10A and 10B to one another, providing each sub-processor system 10 with access to the I/O devices of the other as well as inter-CPU communication. As will be seen, any CPU 12 of the processing system 10 can be given access to the memory of any other CPU 12, although... ..the memory of a CPU 12 by a wayward peripheral device 17.

Preferably, the sub-processor systems 10A/10B are paired as illustrated in Fig. 1A (and Figs 1B and 1C, discussed below), and each sub-processor system 10A/10B pair (i.e., comprising a CPU 12, at least one router 14... ..12A) connects, by a TNet Link L to a router (14A) of the corresponding sub-processor system (e.g., 10A). Conversely, the Y port connects the CPU (12A) to the router (14B) of the companion sub-processor system (10B). This latter connection not only provides a communication path for access by a CPU (12A) to the I/O devices of the other sub-processor system (10B), but also to the CPU (12B) of that system for inter-CPU communication.

Information is communicated between any element of the processing system 10 and any other element (e.g., CPU 12A of sub-processor system 10A) of the system and any other element of the system (e.g., an I/O device associated with an I/O packet interface 16B of sub-processor system 10B) via message "packets." Each message packet is made up of a number of... this reason, a unique method of receiving the symbols at the receiver, using a clock synchronization first-in-first-out (CS FIFO) storage structure (described more fully below), has been developed... operation means just that: the frequencies of the clock signals of the transmitter and receiver units are locked, although not necessarily in phase. Frequency locked clock signals are used to transmit symbols between the routers 14A, 14B and the CPUs 12 of paired sub-processor systems (e.g., sub-processor systems 10A, 10B, Fig. 1A). Since the clocks of the transmitting and receiving element are not phase related, a clock synchronization FIFO is again used -- albeit operating in a slightly different mode from that used for... difference, as will be seen, is due to the fact that pairs of the sub-processor systems 10 can be operated in a synchronized, lock-step mode, called duplex mode, in which each CPU 12 operates to execute the... 1A illustrates another feature of the invention: a cross-link connection between the two sub-processor systems 10A, 10B through the use of additional routers 14 (identified in Fig. 1A as... added routers RX1)), RX2)), RY1)), and RY2)) form a cross-link connection between the sub-processors 10A, 10B (or, as shown, "sides" X and Y, respectively) to couple them to I... shown in Fig. 1A, the routers RX2)) and RY2)) provide the I/O packet interface units 16x and 16y with a dual ported interface. Of course, it will now be evident... lend themselves to being used in a manner that can extend the configuration of the processing system 10 to include additional sub-processor systems such as illustrated in Figs. 1B and 1C. In Fig. 1B, for example, one... of each of the routers 14A and 14B is used to connect the corresponding sub-processor systems 10A and 10B to additional sub-processor systems 10A' and 10B' forming thereby a larger processing system comprising clusters of the basic processing system 10 of Fig. 1.

Similarly, in Fig. 1C the above concept is extended to form an eight sub-processor system cluster, comprising sub-processor systems pairs 10A/10B, 10A'/10B', 10A''/10B'', and 10A'''/10B'''. In turn, each of the sub-processor systems (e.g., sub-processor system 10A) will have essentially the same basic minimum configuration of a CPU 12, a... by a I/O packet interface 16, except that, as Fig. 1C shows, the sub-processor systems 10A and 10B include additional routers 14C and 14D, respectively, in order to extend the cluster beyond sub-processor systems 10A'/10B' to the sub-processor systems 10A''/10B'' and 10A'''/10B'''. As Fig. 1C further illustrates, unused ports 4 and... the routers 14 when configuring the topology of the system 10, any CPU 12 of processing system 10 of Fig. 1C can access any other "end unit" (e.g., a CPU or I/O device) of any of the other sub-processor systems. Two paths are available from any CPU 12 to the last router 14 connecting to the I/O packet interface 16. For example, the CPU 12B of the sub-processor system 10B' can access the I/O 16''' of sub-processor system 10A''' via router 14B (of sub-processor system 10B'), router 14D, and router 14B (of sub-system 10B''') and, via link LA... 10A'''), OR via router 14A (of sub-system 10A'), router 14C, and router 14A (sub-processor system 10A'''). Similarly, CPU 12A of sub-processor system 10A'' may access (via two paths) memory contained in the CPU 12B of sub-processor 10B to read or write data. (Memory accesses by one CPU 12 of another component of the processing system requires, as will be seen, the components seeking access to have authorization to do... prevents corruption of memory data of a CPU by erroneous access.)

The topology of the processing system shown in Fig. 1B is achieved by using port 1 of the routers 14A, 14B, and auxiliary TNet links LA, to connect to the routers 14A', 14B' of sub-processor systems 10A', 10B'. The topology thereby obtained establishes redundant communication paths between any CPU 12 (12A, 12B, 12A', 12B') and any I/O packet interface 16 of the processing system 10 shown in Fig. 1B. For example, the CPU 12A' of the sub-processor system 10A' may access the I/O 16A of sub-processor system 10A by a first path formed by the router 14A' (in port 4, out... shown in Fig. 1B. By interconnecting one port of each router 14 of each sub-processor pair, and using additional auxiliary TNet links LA (illustrated in Fig. 1C with the dotted line connections) between the ports 1 of the routers 14 (14A'' and 14B'') of sub-processor systems 10A'', 10B'' and 10A''', 10B''', two separate, independent data paths can be found between any CPU 12 and any I/O packet interface 16. In this fashion, any end unit (i.e., a CPU 12 or an I/O packet interface 16) will have at least two paths to any other end unit.

Providing alternate paths of access between any two end units (e.g., between a CPU 12 and any other CPU 12, or between any CPU... any two of the remaining fault domains. Here, a fault domain could be a sub-processor system (e.g., 10A). Thus, if the sub-processor system 10A were brought down because of a failure the electrical power being supplied, without... TNet link LA between the routers 14A''' and 14B''', the CPU 12B of the sub-processor system 10B would have lost access to the I/O packet interface 16''' (via router... with the loss of the router 14A (and router 14C) by loss of the sub-processor system 10A, communications between the CPU 12B is still possible via the route of router... equally to CPU 12B. As Fig. 2 shows, the CPU 12A includes a pair of processor units 20a, 20b that are configured for synchronized, lock-step operation in that both processor units 20a, 20b receive and execute identical instructions, and issue identical data and command outputs, at substantially the same moments in time. Each of the processor units 20a and 20b is connected, by a bus 21 (21a, 21b) to a corresponding cache memory 22. The particular type of processor units used could contain sufficient internal cache memory so that the cache memory 22 would not... 22 could be used to supplement any cache memory that may be internal to the processor units 20. In any event, if the cache memory 22 is used, the bus 21 is... 22 address bits, 3 bits of parity covering the address, and 7 control bits.

The processors 20a, 20b are also respectively coupled, via a separate 64-bit address/data bus 23 to X and Y interface units 24a, 24b. If desired, the address/data communicated on each bus 23a, 23b could also be protected by parity, although this will increase the width of the bus. (Preferably, the

processors 20 are constructed to include RISC R4000 type microprocessors, such as are available from the MIPS Division of Silicon Graphics, Inc. of Santa Clara, California.)

The X and Y interface units 24a, 24b operate to communicate data and command signals between the processor units 20a, 20b and a memory system of the CPU 12A, comprising a memory controller (MC... MC halves 26a and 26b) and a dynamic random access memory array 28. The interface units 24 interconnect to each other and to the Mcs 26a, 26b by a 72-bit... accompanied by 8 bits of ECC) are written to the memory 28 by the interface units 24, one interface unit 24 will drive only one word (e.g., the 32 most significant portion) of the doubleword being written while the other interface unit 24 writes the other word of the double word (e.g., the least significant 32-bit portion of the doubleword). In addition, on each write operation the interface units 24a, 24b perform a cross-check operation on the data not written by that interface unit 24 with the data written by the other to check for errors; on read operations... accessed corresponds to the address of the location from which the doubleword was stored.

Interface units 24a, 24b of the CPU 12A form the circuitry to respectively service the X and Y (I/O) ports of the CPU 12A. Thus, the X interface unit 24a connects by the bi-directional TNet Link Lx to a port of the router 14A of the processor system 10A (Fig. 1A) while the Y interface unit 24b similarly connects to the router 14B of the processor system 10B by TNet Link Ly. The X interface unit 24a handles all I/O traffic between the router 14A and the CPU 12A of the sub-processor system 10A. Likewise, the Y interface unit 24b is responsible for all I/O traffic between the CPU 12A and the router 14B of companion sub-processor system 10B.

The TNet Link Lx connecting the X interface unit 24a to the router 14A (Fig. 1) comprises, as above indicated, two 10-bit buses... bus 32x)) carries data incoming from the router 14A. In similar fashion, the Y interface unit 24b is connected to the router 14B (of the sub-processor system 10B) by two 10-bit busses: 30y)) (for outgoing transmissions) and 32y)) (for incoming transmissions), together forming the TNet Link Ly.

The X and Y interface units 24a, 24b are synchronously operated in lock-step, performing substantially the same operations at substantially the same times. Thus, although only the X interface unit 24a actually transmits data onto the bus 30x)), the same output data is being produced by the Y interface unit 24b, and used for error-checking. The Y interface unit 24b output data is coupled to the X interface unit 24a by a cross-link 34y)) where it is received by the X interface unit 24a and compared against the same output data produced by the X interface unit. In this way the outgoing data made available at the X port of the CPU... the port of the CPU 12A is checked. The output data from

the Y interface unit 24b is coupled to the Y port by a 10-bit bus 30y)), and also to the X interface unit 24a by the 9-bit cross-link 34y)) where is checked with that produced by the X interface unit.

As mentioned, the two interface units 24a, 24b operate in synchronous, lock-step with one another, each performing substantially the same... X and/or Y ports of the CPU 12A must be received by both interface units 24a, 24b to maintain the two interface units in this lock-step mode. Thus, data received by one interface unit 24a, 24b is passed to the other, as indicated by the dotted lines and 9... connections 36x)) (communicating incoming data being received at the X port by the X interface unit 24a to the Y interface unit 24b) and 36y)) (communicating data received at the Y port by the Y interface unit 24b to the X interface unit 24a).

Certain more robust operating systems are structured with a fault-tolerant capability in the... example, U.S. Patent No. 4,817,091 teaches a multiprocessor system in which each processor periodically messages each of the processors of the system (including itself), under software control, to thereby provide an indication of continuing operation. Each of the processors, in addition to performing its normal tasks, operates as a backup processor to another of the processors. In the event one of the backup processors fails to receive the messaged indication from a sibling processor, it will take over the operation of that sibling (now thought to be inoperative), in... platform for both types of software. Thus, when a robust operating system is available, the processing system 10 can be configured to operate in a "simplex" mode in which each of... left, in most instances, to software.

Alternatively, for less robust operating systems and software, the processing system 10 provides a hardware-based fault-tolerance by being configured to operate in a... g., CPUs 12A, 12B) are coupled together as shown in Fig. 1A, to operate in synchronized, lock-step fashion, executing the same instructions at the substantially the same moment in time... data and command symbols. In order to simplify the design of the CPU 12, the processors 20 are precluded from communicating directly with any outside entity (e.g., another CPU 12... O device via the I/O packet interface 16). Rather, as will be seen, the processor will construct a data structure in memory and turn over control to the interface units 24. Each interface unit 24 includes a block transfer engine (BTE; Fig. 5) configured to provide a form of... to the destination according to information contained in the message packet.

The design of the processing system 10 permits a memory 28 of a CPU to be read or written by... via the routers 14. Accordingly, before continuing with the description of the construction of the processing system 10, it would be of advantage to understand first the configuration of the data... information.

As indicated, the HADC message packet operates to communicate write data between the end units (e.g., CPU 12) of the processing system 10. Other message packets, however, may be differently constructed because of their function and... CRC. The HC message packet is used to acknowledge a request to write data.

Interface Unit:

The X and Y interface units 24 (i.e., 24a and 24b - Fig. 2) operate to perform three major functions within the CPU 12: to interface the processors 20 to the memory 28; to provide an I/O service that operates transparently to, but under the control of, the processors; and to validate requests for access to the memory 28 from outside sources.

Regarding first the interface function, the X and Y interface units 24a, 24b operate to respectively communicate processors 20a, 20b to the memory controllers (Mcs 26a, 26b) and memory 28 for writing and... fast checking of the data read/written. For example, write operations have the two interface units 24a, 24b cooperating to cross-check the data to be written to ensure its integrity (and at the same time, the interface units 24 will operate) to develop an error correcting code (ECC) that covers, as will be... to have been retrieved from the appropriate address.

With respect to I/O access, the processors 20 are not provided with the ability to communicate directly with the input/output systems... must write data structures to the memory 28 and then pass control to the interface units 24 which perform a direct memory access (DMA) operation to retrieve those data structures, and... indicated in the data structure itself.)

The third function of the X and Y interface units 24, access validation to the memory 28, uses an address validation and translation (AVT) table maintained by the interface units. The AVT table contains an address for each system component (e.g., an I/O... ..the incoming message packets are virtual addresses. These virtual addresses are translated by the interface unit to physical addresses recognizable by the memory control units 26 for accessing the memory 28.

Referring to Fig. 5, illustrated is a simplified block diagram of the X interface unit 24a of the CPU 12A. The companion Y interface unit 24b (as well as the interface units 24 of the CPU 12B, or any other CPU 12) is of substantially identical construction. Accordingly, it will be understood that a description of the interface unit 24a will apply equally to the other interface units 24 of the processing system 10.

As Fig. 5 illustrates, the X interface unit 24a includes a processor interface 60, a memory interface 70, interrupt logic 86, a block transfer engine (BTE) 88, access validation and translation logic 90, a packet transmitter 94, and a packet receiver 96.

Processor Interface:

The processor interface 60 handles the information flow (data and commands) between the processor 20a and the X interface unit 24a. A processor bus 23, including a 64 bit address and data bus (SysAD) 23a and a 9 bit command bus 23b, couples the processor 20a and the processor interface 60 to one another. While the SysAD bus 23a carries memory address and data... ..and qualifying commands carried at substantially the same time on the SysAD bus 23a. The processor interface 60 operates to interpret commands issued by the processor unit 20a in order to pass reads/writes to memory or control registers of the processor interface. In addition, the processor interface 60 contains temporary storage (not shown) for buffering addresses and data for access to... ..26). Data and command information read from memory is similarly buffered en route to the processor unit 20a, and made available when the processor unit is ready to accept it. Further, the processor interface 60 will operate to generate the necessary interrupt signalling for the X interface unit 24a.

The processor interface 60 is connected to a memory interface 70 and to configuration registers 74 by a bi-directional 64 bit processor address/data bus 76. The configuration registers 74 are a symbolic representation of the various control registers contained in other components of the X interface

unit 24a, and will be discussed when those particular components are discussed. However, although not specifically... ..throughout other of the logic that is used to implement the X interface 24a, the processor address/data bus 76 is likewise coupled to read or write to those registers.

Configuration registers 74 are read/write accessible to the processor 20a; they allow the X interface unit to be "personalized." For example, one register identifies the node address of the CPU 12A... ..with the CPU 12A; another, readable only, contains a fixed identification number of the interface unit 24, and still other registers define areas of memory that can be used by, for... ..logic 90, etc.) employing them are discussed.

The memory interface 70 couples the X interface unit 24a to the memory controllers 26 (and to the Y interface unit 24b; see fig. 2) by a bus 25 that includes two 36 bidirectional bit buses 25a, 25b. The memory interface operates to arbitrate between requests for memory access from the processor unit 20, the BTE 88, and the AVT logic 90. In addition to memory accesses from the processor unit 20a, the memory 28 may also be accessed by components of the processing system 10 to, for example, store data requested to be read by the processor unit 20a from an I/O unit 17, or memory 28 may also be accessed for I/O data structures previously set up in memory by the processor unit. Since these accesses are all asynchronous, they must be arbitrated, and the memory interface 70... ..command information accessed from the memory 28 is coupled from the memory interface to the processor interface 60 by a memory read bus 82, as well as to an interrupt logic... ..doubleword quantities. However, while the memory interfaces 70 of both the X and Y interface units 24a and 24b formulate and apply the (64-bit) doubleword to the bus

25, each... ..by the memory interface 70 are coupled to the memory interface by the companion interface unit 24 where they are compared with the same 32 bits for error.

Digressing for the... ..containing interrupt information are received, that information is conveyed to the interrupt logic 86 for processing and posting for action by the processor 20, along with any interrupts generated internal to the CPU 12A. Internally generated interrupts will... ..register 71 (internal to the interrupt logic 86), indicating the cause of the interrupt. The processor 20 can then read and act upon the interrupt. The interrupt logic is discussed more fully below.

The BTE 88 of the X interface unit 24a operates to perform direct memory accesses, and provides the mechanism that allows the processors 20 to access external resources. The BTE 88 can be set-up by the processors 20 to generate I/O requests, transparent to the processors 20 and notify the processors when the requests are complete. The BTE logic 88 is discussed further below.

Requests for... ..8 byte wide format necessary for storing in the memory 28.

Outgoing message packets containing processor originated transaction requests (e.g., a read request asking for a block data from an I/O unit) are monitored by the request transaction logic (RTL) 100. The RTL 100 provides a time... ..will generate an interrupt (handled and reported by the interrupt logic 86) to inform the processor 20 that the request was not honored. In addition, the RTL 100 will validate responses... ..28 (by the DMA operation of the BTE 86) at a location known to the processor 20 so that it can locate the response.

Each of the CPUs 12 are checked... ..discussed. One such check is an on-going monitor of the operation of the interface units 24a, 24b of each CPU. Since the interface units 24a, 24b operate in lock-step synchronism checking can be performed by monitoring the operating states of the paired interface units 24a, 24b by a continuous comparison of certain of their internal states. This approach is implemented by using one stage of a state machine (not shown) contained in the unit 24a of CPU 12A, and comparing each state assumed by that stage with its identical state machine stage in the interface unit 24b. All units of the interface units 24 use state machines to control their operations. Preferably, therefore, a state machine of the memory interface 70 that controls the data transfers between the interface unit 24 and the MC 26 is used. Thus, a selected stage of the state machine used in the memory interface 70 of the interface unit 24a is selected. An identical stage of a state machine of one of the interface unit 24b is also selected. The two selected stages are communicated between the interface units 24a, 24b and received by a compare circuit contained in both interface units 24a, 24b. As the interface units operate lock-step with one another, the state machines will likewise march through the same identical states, assuming each state at substantially the same moments in time. If an interface unit encounters an error, or fails, that activity will cause the interface units to diverge, and the state machines will assume different states. The time will come when... ..that will bring to the attention of the CPUs 12A (or 12B) that the interface units 24a, 24b of that CPU are no longer in lock-step, and to act accordingly... ..X port, receiving only those message packets transmitted by the router 14A of the sub-processor system 10A (Fig. 1A). The Y port is serviced by the Y interface unit 24b to receive message packets from the router 14B of the companion sub-processor system 10B. However, both interfaces (as well as Mcs 26 and processor 20), as has been indicated, are basically mirror images of one another in that both... ..in both structure and function. For this reason, message packet information, received by one interface unit (e.g., 24a) must be passed for processing also to the companion interface unit (e.g., 24b). Further, since both interface units 24a, 24b will assemble the same message packets for transmission from the X or the Y ports, the message packet being transmitted by the interface unit (e.g., 24b) actually being communicated from the associated port (e.g., the Y port) will also be coupled to the other interface unit (e.g., 24a) for cross-checking for errors. These features are illustrated in Figs. 6... ..receiving portions of the packet receivers 96 (96x, 96y) of the X and Y interface units 24a, 24b are broadly illustrated. As shown, each packet receiver 96x, 96y has a clock... ..receive a corresponding one of the TNet Links 32. The CS FIFOs 102 operate to synchronize the incoming command/data symbols to the local clock of the packet receiver 96, buffering... ..104x, coupled to the MUX 104y of the packet receiver 96y of the Y interface unit 24b by the cross-link connection 36x)). In similar fashion, information received at the Y port is coupled to the X interface unit 24a by

the cross-link connection 36y)). In this manner, the command/data symbols of... ..packets received at one of the X, Y ports by the corresponding X, Y, interface unit 24a, 24b is passed to the other so that both will process and communicate the same information on to other components of the interface units 24 and/or memory 28.

Continuing with Fig. 6, depending upon which port X, Y... ..or the other of the CS FIFOs 102x, 102y for communication to the storage and processing logic 110 of the interface unit 24. The information contained in each 9-bit symbol is an 8-bit byte of... ..the encoding of which is discussed below with respect to Fig. 9. The storage and processing logic 110 will first translate the 9-bit symbols to 8-bit data or command... ..the outputs of the CS FIFOs 102x, 102y are also coupled to a command decode unit in addition to the MUX 104. The command decode unit operates to recognize command symbols (differentiating them from data symbols in a manner that is... ..below), decoding them to generate therefrom command signals that are applied to a receiver control unit, a state machine-based element that functions to control packet receiver operations.

As indicated above... ..at the output of the MUX 104, the receiver control portion of the storage control unit enables CRC check logic 106 to calculate a CRC symbol while the data symbols are... ..below, CS FIFOs are found not only in the packet receivers 96 of the interface units 24, but also at each receiving port of the routers 14 and the I/O... ..an even more important part, and perform a unique function, when a pair of sub-processor systems are operating in duplex mode and the two CPUs 12A and 12B of the sub-processor systems 10A, 10B operate in synchronized, lock-step, executing the same instructions at the same time. When operating in this latter... ..difficult to ensure that the clocking regime of the routers 14A and 14B are exactly synchronized to those of the CPUs 12A and 12B - even when using frequency locked clocking. In... ..used to transmit symbols to a CPU 12 and the clock used by an interface unit 24 to receive those symbols.

The structure of the CS FIFO 102 is diagrammatically illustrated... ..i.e., a packet) or IDLE symbols - except during certain situations (e.g., reset, initialization, synchronization and others discussed below). As explained above, each symbol held in the transmit register 120... ..same symbol leaving the storage queue, allowing each symbol entering the storage queue 126 to settle before it is clocked out and passed to the storage and processing units 110x (and 110y) by the MUX 104x (and 104y). Since the transmit and receive clocks... ..functioning in duplex mode) operate to transmit symbols with near frequency clocking. Even so, clock synchronization FIFOs are used at these other ports to receive symbols transmitted with near frequency clocking, and the structure of these clock synchronization FIFOs are substantially the same as that used in frequency locked environments, i.e., that... ..of the storage queue 126 are nine bits wide; in near frequency environments, the clock

synchronization FIFOs use symbol locations of the queue 126 that are 10 bits wide, the extra... ..the faster clock source. To handle this clock drift, the two pointers are effectively re-synchronized periodically.

When the CPUs 12 are paired and operating in duplex mode, all four interface units 24 operate in lock-step to, among other things, transmit the same data and receive... ..simplex mode, each independent of the other, clocking need only be near frequency.

The interface unit 24 receives a SYNC CLK signal that is used in combination with a SYNC command symbol to initialize and synchronize the Rcv register 124 to the transmitting router 14. When using either near frequency or... ..102X preferably begin from some known state. Incoming symbols are examined by the storage and processing units 110 of the packet receivers 96. The storage and processing units look for, and act upon as appropriate, command symbols. Pertinent here is that when the... ..receives a SYNC command symbol it will be decoded and detected by the storage and processing unit 110. Detection of the SYNC command symbol by the storage and processing unit 110 causes assertion of a RESET signal. The RESET signal, under synchronous control of the SYNC CLK signal, is used to reset the input buffers (including the clock synchronization buffers) to predetermined states, and synchronize them to the routers 14.

The synchronization of the CS FIFOs 102 of the interface units 24 those of one or both routers 14A, 14B is discussed more fully below in the section discussing synchronization.

Packet Transmitter:

Each interface unit 24 is assigned to transmit from and receive at only one of the X or Y ports of the CPU 12. When one of the interface units 24 transmits, the other operates to check the data being transmitted. This is an important... ..shows, in abbreviated form, the packet transmitters 94x, 94y of the X and Y interface units 24a, 24b, respectively. Both packet transmitters are identically constructed, so that discussion of one (packet... ..logic 152 that receives, from the BTE 88 or AVT 90 of the associated interface unit (here, the X interface unit 24a) the data to be transmitted - in doubleword (64-bit) format. The packet assembly logic... ..and Y ports: they are either symbols that make up a message packet in the process of being transmitted, or IDLE symbols, or other command symbols used to perform control functions... ..154, 156. The output of the multiplexer 154 connects to the X port. (The interface unit 24b connects the output of the multiplexer 154 to the Y port.) The multiplexer 156... ..link 34x)) to the checker logic 160 of the packet transmitter 94y (of the interface unit 24b).

A selection (S) input of the multiplexers receives a 1-bit output from an... ..is accessible to the MP 18 via an OLAP (not shown) formed in the interface unit 24, and is written with information that "personalizes," among other things, the interface units 24 Here, the X/Y stage of the configuration register 162 configures the packet transmitter 94x of the X interface unit 24a to communicate the X encoder 150x output to the X port; the output of... ..traffic is present, the operation of the two packet interfaces 94 (and, thereby, the interface units 24 with which they are associated) are continually monitored. Should one of the checkers detect... ..will be asserted, resulting in an internal interrupt being posted for appropriate action by the processors 20.

Message packet traffic operates in the same manner. Assume, for the moment, that the... ..that information, a byte at a time, to the X encoder 150x of both interface units 96, which will translate each byte to encoded 9-bit form. The output of the... ..is checked with that from the packet transmitter 94x. Again, the operation of the interface units 24a, 24b, and the packet transmitters they contain, are inspected for error.

In the same... ..monitored.

Returning for the moment to Fig. 5, if the outgoing message packet is a processor initiated transaction (e.g., a read request), the processors 20 will expect a message packet to be returned in response. Thus, when the BTE... ..will issue a timeout signal to the interrupt logic (Fig. 14A) to thereby notify the processors 20 of the absence of a response to a particular transaction (e.g., a read... ..the access, to name just a few. Also, the area of memory of the memory unit 28 desired to be accessed are identified in the message packets by virtual or I... ..virtual addresses be translated to physical addresses of the memory 28. Finally, interrupts generated by units or elements external to the CPU 12A, are transmitted via message packets to interrupt the processors 20, which are also written to memory 28 when received. All this is handled by the interrupt logic and AVT logic 86, 90.

The AVT logic unit 90 utilizes a table (maintained by the processor 20 in memory 28) containing AVT entries for each possible external source permitted access to the memory 28. Each AVT entry identifies a specific source element or unit and the particular page (a page being nominally 4K (4096) bytes), or portion of a... ..expected" memory accesses. Expected memory accesses are those initiated by the CPU 12 (i.e., processors 20) such as a read request for information from an I/O device. These latter memory accesses are handled by a transaction sequence number (TSN) assigned to each processor initiated request. At about the time the read request is generated, the processors 20 will allocate an area of memory for the data expected to be received in... ..and 26b are, in turn, respectively coupled to the memory interfaces 70 of each interface unit 24a, 24b. The 64-bit doublewords are written to the memory 28 with the upper... ..check bits respectively from the memory interfaces 70 (70a, 70b) of each of the interface units 24a, 24b (Fig. 5).

Referring to Fig. 10, each memory interface 70 receives, from either the bus 82 from the processor interface 60 or the bus 83 from AVT logic 90 (see Fig. 5), of the associated interface unit 24, 64 bits of data to be written to memory. The busses 76 and 83... other for cross-checking between them. Thus, for example, the memory interface 70a (of interface unit 24a) will drive the MC 26a with the "upper" 32 bits of the 64 bits... are check bits, leaving 40 bits unused.

Access Validation:

As previously indicated, components of the processing system 10 external to the CPU 12A (e.g., devices of the I/O packet... not without qualification. Access validation, as implemented by the AVT logic 90 of the interface units 24, operates to prevent the content of the memory 28 from being corrupted by erroneously... Accesses to the memory 28 are validated by the AVT logic 90 of each interface unit 24 (Fig. 5), using all of six checks: (1) that the CRC of the message... also are permitted the particular message packet source.

The access validation mechanism of the interface unit 24a, AVT logic 88, is shown in greater detail in Fig. 11. Incoming message packets... and post an interrupt to the interrupt logic 86 (Fig. 5) for action by the processor 20.

The mask operation permits the size of the table of AVT entries to be varied. The content- of the AVT mask register 175 is accessible to the processor 20, permitting the processors 20 to optionally select the size of the AVT entry table. A maximum AVT table... 172 allows the AVT size to be matched to the needs of the system. A processing system 10 that includes a larger number of external elements (e.g., the number of... amount of the memory space of memory 28 to the AVT entries. Conversely, a smaller processing system 10, with a smaller number of external elements will not have such a large... set to a logic "ZERO" indicate an nonexistent TNet address, outside the limits of the processing system 10. A received packet with a TNet address outside the allowable TNet range will... in Fig. 11 as being held in the AVT entry register 180 during the validation process. AVT entries have two basic formats: normal and interrupt. The format of a normal AVT... of the AVT input register 170) will result in an error being posted to the processor via an interrupt.

A 12-bit "Permissions" field is included in t AVT...path=0). Denials are logged as interrupts with the interrupt logic, and reported to the processor 20 - if the E field is set to a state ("ONE") that enables error-reporting... e.g., to a "ONE"), the other fields (Upper Bound, etc.) gain new definitions for processing interrupt writes and managing interrupt queues. This is discussed in more detail below in connection... memory 28 will be handled. Set to one state, the requested write operation will be processed normally; set to a second state, write requests specifying addresses with a fractional cache line... be written to a specific queue (interrupt queue) in memory 28, with signalling provided the processors 20 to indicate that an interrupt has been received and "posted," and ready for servicing by the processors 20. Since the interrupt queues are at specific memory locations, the processor can obtain the interrupt data when needed.

An AVT interrupt entry for an interrupt may... by the interrupt logic 86, and extracted from the head of the queue by the processor 20 when servicing the interrupt.

The AVT interrupt entry also includes a 20-bit segment ("Source ID") containing source ID information, identifying the external unit seeking attention by the interrupt process. If the source ID information of the AVT interrupt entry does not match that contained... class" of the interrupt that is used to determine the interrupt level set in the processor 20 (described more fully below); (2) a queue number that is used to select, as... capability to deliver interrupts to a CPU 12 for servicing. For example, an I/O unit may be unable to complete a read or write transaction issued by a CPU because... identify the recipient. These and other errors, exceptions, and irregularities, noted by the I/O

units, or the I/O Interface elements, can become the a condition that requires the intervention... the AVT entry register 180 for use by the interrupt logic 86 of the interface unit 24 (Fig. 5), illustrated in greater detail in Fig. 14A.

It is interrupt logic 86... four circular queues specified by the base address information contained in the AVT entry. The processor(s) 20 will then be notified, and it will be up to them as to... selected tail queue register 256 by combiner circuit 270, the output of which is the processed by the "mod z" circuit 273 to turn new offset into the queue at which... signal. The Queue Full warning signal becomes an "intrinsic" interrupt that is conveyed to the processor units 20 as a warning that if the matter is not promptly handled, later-received interrupt... will be discarded.

Incoming message packet interrupts will cause interrupts to be posted to the processor 20 by first setting one of a number of bit positions of an interrupt register 280. Multi-entry queued interrupts are set in interrupt registers 280a for posting to the processor 20; single-entry queue interrupts use interrupt register 280b. Which bit is set depends upon... multi-entry queued interrupts, soon after a multi-entry queued interrupt is determined, the interface unit will assert a corresponding interrupt signal (I1) that is applied to decode circuit 283. Decode... of register 280a to set, thereby providing advance information concerning the received interrupt to the processor(s) 20, i.e., (1) the type of interrupt posted, and (2) the class of... to one another by a compare circuit 279. The update register is writable by the processor 20 to select a register pair for comparison. If the content of the two selected... cleared.

Digressing for the moment, there are two basic types of interrupts that concern the processors 20: those interrupts that are communicated to the CPU 12 by message packets, and those... the seven interrupt postings to a latch 288, from which they are coupled to the processor 20 (20a,20b) which has an interrupt register for receiving holding the postings.

In addition... change in interrupts (either an interrupt has been serviced, and its posting deleted by the processor 20, or a new interrupt has been posted), a "CHANGE" signal will be issued to the processor interface 60 to inform it that an interrupt posting change has occurred, and that it should communicate the change to the processor 20.

Preferably, the AVT entry register 180 is configured to operate like a single line... such as set-associative, fully-associate, or direct-mapped, to name a few.

Coherency:

Data processing systems that use cache memory have long recognized the problem of coherency: making sure that... the incoming packet is permitted access are applied to a boundary crossing (Bdry Xing) check unit 219. Boundary check unit 219 also receives an indication of the size of the cache block the CPU 12... Len field of the header information from the AVT input register 170. The Bdry Xing unit determines if the data of the incoming packet is not aligned on a cache boundary... time an interrupt will be written to the queued interrupt register 280, to alert the processors 20 that a portion of the incoming data is located in the special queue.

In... not, the packet (both header and data) is written to a special queue, and the processors so notified by the intrinsic interrupt process described above. The processors may then move the data from the special queue to cache 22, and later write... the cache 22 and the memory 28 is preserved.

Block Transfer Engine (BTE):

Since the processor 20 is inhibited from directly communicating (i.e., sending) information to elements external to the... indirect method of information transmission.

The BTE 88 is the mechanism used to implement all processor initiated I/O traffic to transfer blocks of information. The BTE 88 allows creation of... BTE registers 300, 302 whose content is coupled to the MUX 306 (of the interface unit 24a; Fig. 5) and used to access the system memory 28 via the memory controllers... BTE data structure 304 in the memory 28 of the CPU 12A (Fig. 2). The processors 20 will write a data structure 304 to the memory 28 each time information is... begin on a quadword boundary, and the BTE registers 300, 302 are writable by the processors 20 only. When a processor does write one of the BTE registers 300, 302, it does so with a word...

...the request bit (rc0, rc1) to a clear state, which operates to initiate the BTE process, which is controlled by the BTE state machine 307.

The BTE registers 300, 302 also... ..cause (ec) bit differentiates time-outs and NAKs.

When information is being transferred by the processors 20 to an external unit, the data buffer portion 304b of the data structure 304 holds the information. When information from an external unit is received by the processors 20, the data buffer portion 304b is the location targeted to hold the read response information.

The beginning of the data structure 304, portion 304a written by the processor 20, includes an information field (Dest), identifying the external element which will receive the packet... ..the transmitted data is to be written. This information is used by the packet transmitter unit 120 (Fig. 5) to assemble the packet in the form shown in Figs. 3-4... ..list (el) bit, when set, indicates the end of the chain, and halts the BTE processing.

The interrupt completion (ic) bit, when set, will cause the interface unit 24a to assert an interrupt (BTECmp) which sets a bit in the interrupt register 280... ..the chain pointer).

The interrupt time-out (it) bit, when set, will cause the interface unit 24a to assert an interrupt signal for the processor 20 if the acknowledgement of the access times-out (i.e., if the request timer... ..time), or elicits a NAK response (indicating that the target of the request could not process the request).

Finally, if the check sum (cs) bit is set, the data to be... ..containing the data from which the check sum was formed.

To sum up, when the processors 20 of the CPU 12A desire to send data to an external unit, they will write a data structure 304 to the memory 28, comprising identifier information in portion 304a of the data structure, and the data in the buffer portion 304b. The processors 20 will then determine the priority of the data and will write the BTE register... ..information, and sent.

If the data structure 304 indicates a read request (i.e., the processors 20 are seeking data from an external unit - either an I/O device or a CPU 12), the Len and Local Buffer Ptr... ..receiver 100 (Fig. 5) until the local memory write operation is executed.

Responses to a processor-generated read request to an external unit are not processed by the AVT table logic 146. Rather, when the processors 20 set up the BTE data structure, a transaction sequence number (TSN) is assigned the... ..the BTE 88, which will be an HAC type packet (Fig. 4) discussed above. The processors 20 will also include an memory address in the BTE data structure at which the... ..302, assume that the foregoing transfer of data from the CPU 12A to an external unit is of a large block of information. Accordingly, a number of data structures would be set up in memory 28 by the processors 20, each (except the last) including a chain pointer to additional data structures, the sum... ..sent. Assume now that a higher priority request is desired to be made by the processors 20. In such a case, the associated data structure 304 for such higher priority request... ..with another BTE operation descriptor.

Memory Controller:

Returning, for the moment, to Fig. 2, interface units 24a, 24b access the memory 28 via a pair of memory controllers (MC) 26a, 26b. The Mcs provide a fail-fast interface between the interface units 24 and the memory 28. The Mcs 26 provide the control logic necessary for accessing... ..in dynamic random access memory (DRAM) logic). The Mcs receive memory requests from the interface units 24, and execute reads and writes as well as providing refresh signals to the DRAMs... ..to provide a 72 bit data path between the memory array 28 and the interface units 24a, 24b, which utilize an SBC-DBD-SbD ECC scheme, where b=4, on a... ..26a, 26b to work together and simultaneously supply a 64-bit word to the interface units 24 with minimum latency, one-half of which (D0) comes from the MC 26a, and the other half (D1) comes from the other MC 26b. The interface unit 24 generate and check the ECC check bits. The ECC scheme used will not only... ..26 bus 25, as well as in internal registers.

From the viewpoint of the interface units 24, the memory 28 is accessed with two instructions: a "read N doubleword" and a... ..doubleword read or a block read format. The signal called "data valid" tells the interface units 24 two cycles ahead of time that read data is being returned or not being returned.

As indicated above, the maintenance processor (MP 18; Fig. 1A) has two means of access to the CPUs 12. One is... ..18 will write a register contained in the OLAP 285 with instructions that permit the processors 20 to build an image of a sequence of instructions in the memory that will permit them (the processors 20) to commence operation, going to I/O for example to transfer instructions and data from an external (storage) device that will complete the boot process.

The OLAP 285 is also used by the processors 20 to communicate to the MP 18 error indications. For example, if one of the interface units 24 detect a parity error in data received from the memory controller 26, it will... ..and address transfers on the bus 25 between the MC 26a and the corresponding interface unit 24a. The addressing and data transfers on the DRAM data bus, as well as generation... ..the CPU 12.

Packet Routing:

The message packets communicated between the various elements of the

processing system 10 (e.g., CPUs 12A, 12B, and devices coupled to the I/O packet... ..First, each TNet Link L connects to an element (e.g., router 14A) of the processing system 10 via a port that has both receive and transmit capability. Each transmit port... ..cycle (i.e., each clock period) of the T(underscore)Clk so that the clock synchronization FIFO at the receiving end of the transmission will maintain synchronization.

Clock synchronization is dependent upon the mode in which the processing system 10 is operated. If operating in the simplex mode in which the CPUs 12A... ..connect directly to the CPUs may drift with respect to each other. Conversely, when the processing system 10 operates in a duplex mode (e.g., the CPUs operate in synchronized, lock-step operation), the clocks between routers 14 and the CPUs 12 to which they... ..not necessarily phase-locked).

The flow of data packets between the various elements of the processing system 10 is controlled by command symbols, which may appear at any time, even within... ..initiated by a CPU 12, or MP 18, and promulgated to all elements of the processing system 10 by the routers 14 to communicate an event requiring software action by all... ..command symbol is used in conjunction with near frequency operation as an aid to maintaining synchronization between the two clock signals that (1) transfer each symbol to, and load it in each receiving clock synchronization FIFO, and (2) that retrieves symbols from the FIFO.

SLEEP: This command symbol is sent by any element of the processing system 10 to indicate that no additional packet (after the one currently being transmitted, if... ..received.

SOFT RESET (SRST): The SRST command symbol is used as a trigger during the processes ("synchronization" and "reintegration," described below) that are used to synchronize symbol transfers between the CPUs 12 and the routers 14A, 14B, and then to place... ..SYNC command symbol is sent by a router 14 to the CPU 12 of the processing system 10 (i.e., the sub-processor systems 10A/10B) to establish frequency-lock synchronization between CPUs 12 and routers 14A, 14B prior to entering duplex mode, or when in duplex mode to request synchronization, as will be discussed more fully below. The SYNC command symbol is used in conjunction... ..or duplex to simplex), among other things, as discussed further below in the section on Synchronization and Reintegration.

THIS LINK BAD (TLB): When any system element receiving a symbol from a TNet link L (e.g., a router, a CPU, or an I/O unit) notes an error when receiving a command symbol or packet, it will send a TLB... ..identical pairs of symbols that are compared to one another when pulled from the clock synchronization FIFOs..The DVRG command symbol signals the CPU 12 that a mis-compare has been noted. When received by the CPUs, a divergence

detection process is entered whereby a determination is made by the CPUs which CPU may be failing... ..command symbols described above operate to control message flow between the various elements of the processing system 10 (e.g., CPUs 12, router 14, and the like), using principally the BUSY... ..particular TNet port however, an "end node" (i.e., a CPU 12 or I/O unit 17 - Fig. 1) may not assert backpressure because one of its transmit ports is backpressured... ..Improperly addressed packets are discarded by the router 14.

When a system element of the processing system 10 receives a BUSY command symbol on a TNet link L on which it... ..other command symbols (READY, BUSY, etc.).

Whenever a TNet port of an element of the processing system 10 detects receipt of a READY command symbol, it will terminate transmission of FILL... ..receives.

As will be seen, all elements (e.g., router 14, CPUs 12) of the processing system 10 that connect to a TNet link L for receiving transmitted symbols will receive those symbols via a clock synchronization (CS) FIFO. For example, as discussed above, the interface units 24 of CPUs 12 include all CS FIFOs 102x, 102y (illustrated in Fig. 6). The... ..depth to allow for speed matching, and the elastic FIFOs must provide sufficient depth for processing delays that may occur between transmission of a BUSY command symbol during receipt of a... ..another data byte in packet B. As packet A progresses to the next router, the process would be repeated. If the router 14 displaces more data bytes than the FIFO can... ..irrespective of its own findings.

SLEEP Protocol:

The SLEEP protocol is initiated by a maintenance processor via a maintenance interface (an on-line access port - OLAP), described below. The SLEEP protocol... ..reintegrate a slice of the system 10. Routers 14 must be idle (no packets in process) in order to change modes without causing data loss or corruption. When a SLEEP command symbol is received, the receiving element of processing system 10 inhibits initiation of transmission of any new packet on the associated transmit port...The HALT command symbol provides a mechanism for quickly informing all CPUs 12 in a processing system 10 that is necessary to terminate I/O activity (i.e., message transmissions between... ..CPUs that receive HALT command symbols on either of their receive ports (of the interface units 24) will post an interrupt to the interrupt register 280 if the system halt interrupt... ..interrupt; Fig. 14A).

The CPUs 12 may be provided with the ability to disable HALT processing. Thus, for example, the configuration registers 75 of the interface units 24 can include a "halt enable register" that, when set to a predetermined state (eg., ZERO) disables HALT processing, but reporting detection of a HALT symbol as an error.

Router Architecture:

Referring now to... ..simplified block diagram of the router 14A is illustrated. The other routers 14 of the processing system 10 (e.g., routers 14B, 14', etc.) are of substantially identical construction and, therefore... ..these ports 4, 5 are structured to operate in a frequency locked environment when a processing system 10 is set for duplex mode operation. In addition, when in duplex mode, a... ..5025)) will receive the command/data symbols from the CPUs, pass them through the clock synchronization FIFOs 518 (discussed further below), and compare each symbol exiting the clock synchronization FIFOs with a gated compare circuit 517. When duplex operation is entered, a configuration register... ..517 to activate the symbol by symbol comparison of the symbols emanating from the two synchronization FIFOs 518 of the router input logic 502 for the ports 4 and 5. Of... ..to that received, at substantially the same time, by the other port input.

To maintain synchronization in the duplex mode, the two port outputs of the router 14A that transmit to... ..mode, are duplicated by the routers 14, and returned to both CPUs.) The output logic units 5044)), 5045)) that are coupled directly to the CPUs 12 will both receive symbols from... ..message packet identifies only one of the duplexed CPUs 12, e.g., CPU 12A) in synchronized fashion, presenting those symbols in substantially simultaneous fashion to the two CPUs 12. Of course, the CPUs 12 (more accurately, the associated interface units 24) receive the transmitted symbols with synchronizing FIFOs of substantially the same structure as that illustrated in Fig. 7A so that, even...

...from the FIFO structures by both CPUs 12 on the same instruction cycle, maintaining the synchronized, lock-step operation of the CPUs 12 required by the duplex operating mode.

As will... ..in conjunction with configuration data written to registers contained in control logic 509 by the maintenance processor 18 (via the on-line access port 285' and serial bus 19A; see Fig. 1A... ..links L. The input logic 505 of each port input 502 also assists in maintaining synchronization - at least for those ports sending symbols in the near-frequency environment - by removing received... ..slower-receiving element receiving symbols from a faster-sending element could overload the input clock synchronization FIFO of the slower-receiving element. That is, if a slower clock is used to pull symbols from the clock synchronization FIFO put there by a faster clock, ultimately the clock synchronization FIFO will overflow.

The preferred technique employed here is to periodically insert SKIP symbols in... ..stream to avoid, or at least minimize, the possibility of an overflow of the clock synchronization FIFO (i.e., clock synchronization FIFO 518; Fig. 20A) of a router 14 (or CPU 12) due to a T... ..being slightly higher in frequency than the local clock used to pull symbols from the synchronization FIFO. Using SKIP symbols to by-pass a push (onto the FIFO) operation has the... ..stall each time a SKIP command symbol is received so that, insofar as the clock synchronization FIFO is concerned, the transmitting clock that accompanied the SKIP symbol was missing.

Thus, logic... ..the port inputs 502 will recognize, and key off receipt of, SKIP command symbols for synchronization in the near frequency clocking environment so that nothing is pushed onto the FIFO, but... ..14, or between routers 14, or between a router 14 and an I/O interface unit 16A - Fig. 1) at a 50 Mhz rate, this allows for a worst case frequency... ..symbol by supplying FILL or IDLE symbols (which are received and pushed onto the clock synchronization FIFOs, but are not passed to the elastic FIFOs). In short, each elastic FIFO 506... ..received symbols are then communicated from the input register 516 and applied to a clock synchronization FIFO 518, also by the T(underscore)Clk. The clock synchronization FIFO 518 is logically the same as that illustrated in Figs. 8A and 8B, used in the interface units 24 of the CPUs 12. Here, as Fig. 20A shows, the clock synchronization FIFO 518 comprises a plurality of registers 520 that receive, in parallel, the output of... ..516. Associated with each of the registers 520 is a two-stage validity (V) bit

synchronizer 522, shown in greater detail in Fig. 20B, and discussed below. The content of each... ..registers 520, together with the one-bit content of each associated two-stage validity bit synchronizer 522, are applied to a multiplexer 524, and the selected register/synchronizer pulled from the FIFO, and coupled to the elastic FIFO 506 by a pair of... ..is determined the state of the Push Select signal provided by a push pointer logic unit 530; and, selection of which register 520 will supply its content, via the MUX 524... ..and loading of the register 520 selected by the push pointer logic 530. Similarly, the synchronization FIFO control logic 534 receives the clock signal local to the router (Rcv Clk) to... ..pointer logic 532.

Digressing for a moment, and referring to Fig. 20B, the validity bit synchronizer 522 is shown in greater detail as including a D-type flip-flop 541 with... ..530 (Fig. 20A) selects the register 520 of the FIFO with which the validity bit synchronizer is associated for receipt of the next symbol - if not a SKIP symbol.

The delay... ..Truth Table, below). The D-type flip-flop 543 acts as an additional stage of synchronization, ensuring a stable level at the V output relative to the local Rec Clk. The... ..flip-flop 542, allowing the Pull signal (a periodic pulse from the sync FIFO Control unit 534) to clear the validity bit on this validity synchronizer 522 when the associated register 520 has been read.

In summary, the validity synchronizer 522 operates to assert a "valid" (V) signal when a symbol is loaded in a... ..blocked from being routed out a particular port because another message is already in the process of being routed out that port. However, that other message in turn is also blocked... ..an incoming message packet bound for the CPUs will be replicated by the crossbar logic unit by routing the message packet to both port output 5044)) and

5045)) at the same... ..P) identifies which of path (X or Y) should be used for accessing two sub-processing the device.

The routers 14 provide a capability of constructing a large, versatile routing network for, for example, massively parallel processing architectures. Routers are configured according to their location (i.e., level) in the network by...expansion registers 509j)) and 509k)) are such that bits "def" are used in the algorithmic process, then bits "abc" of the Region ID are compared to the content of the Device... ..content of the route to default register 509f))) to the final stage of the selection process: check logic 602. Check logic 602 operates to check the status of the port output... ..a lower level router, and may be located in one or another of the sub-processing systems 10A, 10B. Whether a router is an upper level or lower level router depends... ..of CPUs 12 and I/O devices 16 to one another, forming a massively parallel processing (MPP) system. Other such MPP systems may exist, and it is those routers configured as... ..captured. As soon as the message packet's Destination ID is so captured, the selection process begins, proceeding to the development of a target port address that will be used to... ..an error that will be posted to the MP18 via the router's (or interface unit's) OLAP for action.

Digressing, it should be appreciated that these protocol rules observed by the routers 14 are also observed by the CPUs 12 (i.e., interface units 24) and I/O packet interfaces 17.

Finally, when the router 14A is in the... ..directly with the CPUs 12A, 12B, and duplex mode is used, a duplex operation logic unit 638 is utilized to coordinate the port output connected to one of the CPUs 12A... ..was able to write instructions to the OLAP 285 that would be executed by the processors 20 to build a small memory image and routine to permit the CPU 12 to... ..the clock generation circuit design. There will be one clock generator circuit in each sub-processor system 10A/10B (Fig. 1) to maintain synchronism. Designated generally with the reference numeral 650... ..used by the various elements (e.g. CPU. 12, routers 14, etc.) of the sub-processor system containing the clock circuit 650 (e.g., 10A).

The clock generator 654 is shown... ..The 50 Mhz clock signals produced by the counter 663 are distributed throughout the sub-processor system where needed.

Turning now to Fig. 25, there is illustrated the interconnection and use... ..the clock circuits 650 used to develop synchronous clock signals for a pair of sub-processor systems 10A, 10B (Fig. 1) for frequency locked operation. As illustrated in Fig. 25, the two CPUs 12A and 12B of the sub-processor systems 10A, 10B each have a clock circuit 650, shown in Fig. 25 as clock... ..654B of both CPUs 12. A driver and signal line 667 interconnects the two sub-processor systems to deliver the M(underscore)CLK signal developed by the oscillator circuit 652A to the clock generator 654B of the sub-processor system 10B. For fault isolation, and to maintain signal quality, the M(underscore)CLK signal is delivered to the clock generator 654A of the sub-processor system 10A through a separate driver and a loopback connection 668. The reason for the... ..the cable (not shown) will establish the connection shown in Fig. 25 between the sub-processor systems 10A, 10B; connected another way, the connections will be similar, but the oscillator 652B... ..Fig. 25, the M(underscore)CLK signal produced by the oscillator circuit 652A of sub-processing system 10A is used by both sub-processing systems 10A, 10B as their respective SYNC CLK signals and the various other clock signals... ..produced by the clock generators 654A, 654B. Thereby, the clock signals of the paired sub-processing systems 10A, 10B are synchronized for the frequency locked operation necessary for duplex mode.

The VCXOs 662 of the clock... ..This allows both clock generators 654A, 654B to continue to provide to the two sub-processing systems 10A, 10B clock signals in the face of improper operation of the oscillator circuit 652A, although the sub-processor systems may no longer be frequency-locked.

The LOCK signals asserted by the phase comparators... ..LOCK signal signifies that the 50 Mhz signals produced by a clock generator 654 are synchronized, both in phase and in frequency, to the M(underscore)CLK signal. Thus, if either... ..signal that accompanies the symbol stream, and is used to push symbols onto the clock synchronizing FIFO of the receiving element (router 14, or CPU 12) is substantially identical in frequency... ..not phase, to that of

the receiving element used to pull symbols from the clock synchronization FIFOs. For example, referring to Fig. 23, which illustrates symbols being sent from the router... clock (Local Clk). The former (Rcv Clk) is used to push symbols onto the clock synchronization FIFOs 126 of each CPU, whereas the latter is used to pull symbols from the... much higher frequency clock signal. In such situations provision must be made to ensure that synchronization is maintained between the two CPUs as to symbols pulled from the clock synchronization FIFOs 126 of each.

Here, a constant ratio clocking mechanism is used to control operation of the two clock synchronization FIFOs 126, providing the clock signal that pulls symbols from the two FIFOs at the... control mechanism is shown, designated with the reference numeral 70. As Fig. 26A illustrates, clock synchronization FIFO control mechanism 700 includes an pre-settable, multi-stage serial shift register 702, the... ratio of the clock signal at which symbols are communicated and pushed onto the clock synchronization FIFOs 126 to the frequency of the clock signal used locally. Here, a 15 stages... that will be used as the Local Clk signal to pull symbols from the clock synchronization FIFOs 126, and to operate (update) the pull pointer counter 130. The selected output is... of the CPU 12 to the clock signal used to push symbols onto the clock synchronization FIFO 126, Rcv Clk, the serial shift register is preset so that M stages of... duplexed CPUs 12 with a 50 Mhz clock. Thus, symbols are pushed onto the clock synchronization FIFOs 126 of the CPUs at a 50 Mhz rate. Assume further that the clock... of the MUX 704, which produces the clock signal that pulls symbols from the clock synchronization FIFOs 126, Rcv Clk, will contain, for each 100 ns period, five clock pulses. Thus... five symbols will be pushed onto, and five symbols will be pulled from, the clock synchronization FIFOs 126.

This example is symbolically shown in Fig. 26B, while the timing diagram shown... labelled "IN" in Fig. 27) of the Rcv Clk will push symbols onto the clock synchronization FIFOs 126. During that same 100 ns period, the serial shift register 702 circulates a... clocks which would require additional storage (i.e., an increase in the size of the synchronization FIFO) and impose more latency.

The constant ratio clock circuit presented here (Figs. 26) is... frequency to a clock regime of a different, higher frequency. The use of a clock synchronization FIFO is necessary here for compensating effects of signal delays when operating in synchronized, duplexed mode to receive pairs of identical command/data symbols from two different sources. However synchronization FIFO. Transferring data ...a wide range of possible clock ratios.

I/O Packet Interface:

Each of the sub-processor systems 10A, 10B, etc. will have some input/output capability, implemented with various peripheral units, although it is conceivable that the I/O of other sub-processor systems would be available so that a sub-processing system may not necessarily have local I/O. In any event, if local I/O... device (e.g., a signal line) would be received by the I/O packet interface unit 16 and used to form an interrupt packet that is sent to the CPU 12... OLAP bus, configuration information.

On-Line Access Port:

The MP 18 connects to the interface

unit 24, memory controller (MC) 26, routers 14, and I/O packet interfaces with interface signals... OLAP 258 is essentially the same, regardless of what element (e.g. router 14, interface unit 24, etc.) it is used with. Fig. 28 diagrammatically illustrates the general structure of the... circuit chip used to implement certain of the elements discussed herein. For example, each interface unit 24, memory controller 26, and router 14 is implemented by an application specific integrated circuit... of the OLAP 158 shown in Fig. 28 describes the OLAP associated with the interface unit 24, the MC 26, and the router 14 of the system.

As Fig. 28 shows... ..asymmetric variables, a "soft-vote" (SV) logic element 900 (Fig. 30A) is provided each interface unit 24 of each CPU 12. As Fig. 30 illustrates, the SV logic elements 900 of each interface unit 24 are connected to one another by a 2-bit SV bus 902, comprising bus lines 902a and 902b. Bus lines 902a carry one-bit values from the interface units 24 of CPU 12A to those of CPU 12B. Conversely, bus line 902b carries one... ..the CPU 12A.

Illustrated in Fig. 30B, is the SV logic element 900a of interface unit 24a of CPU 12A. Each SV logic element 900 is substantially identical in construction and... ..900a should be understood as applying equally to the other logic elements 900a (of interface unit 24b, CPU 12A), and 900b (of the interface units 24a, 24b of CPU 12B) unless noted otherwise. As Fig. 30B illustrates, the SV logic... ..the logic elements 900a (as well as its own). In this manner the two interface units 24a, 24b of the CPU 12A can communicate asymmetrical variables to each other.

In a... ..to the remote register 907 of logic element 902a (and that of the other interface unit 24b).

The logic elements 902 form a part of the configuration registers 74 (Fig. 5). Thus, they may be written by the processor unit(s) 20 by communicating the necessary data/address information over at least a portion of... ..local and remote registers 906 and 907.

The MUX 914 operates to provide each interface unit 24 of CPU 12A with selective use of the bus line 902a for the SV logic elements 900a, or for communicating a BUS ERROR signal if encountered during the reintegration process (described below) used to bring a pair of CPUs 12 into lock-step, duplex operation... ..same time, write the enable registers 912 of the logic element 900 of both interface units 24 of each CPU. One of the two logic elements 900 of each CPU will... ..it is the output enable registers 912 associated with the logic elements 900 of interface units 24a of both CPUs 12A, 12B that are written to enable the associated drivers 916. Thus, the output registers 904 of the interface units 24a of each CPU will be communicated to the bus lines 902; that is, the... ..to the bus line 902a, while the output register associated with logic element 900b, interface unit 24a of CPU 12B is communicated to bus line 902b. The CPUs 12 will both... ..again written by each CPU, followed again by reading the remote input registers 907. This process is repeated, one bit at a time, until the entire variable is communicated from the... ..each CPU 12 to the remote input register of the other. Note that both interface units 24 of CPU 12B will receive the bit of asymmetric information.

One example of use... ..elements 900 are also used to communicate bus errors that may occur during the reintegration process to be described. When reintegration is being conducted, a REINT signal will be asserted. As... ..ERROR signal is selected by the MUX 914 and communicated to the bus line 902a.

Synchronization:

Proper operation of the sub-processing systems 10A, 10B (Figs. 1A, 2) whether operating independently (simplex mode), or paired and operating in synchronized lock-step (duplex mode), requires assurance that data communicated between the CPUs 12A, 12B and... ..routers 14A, 14B will be received properly, and that any initial content of the clock synchronization FIFOs 102 (of CPUs 12A, 12B; Fig. 5) and 519 (of routers 14A, 14B; Fig... ..erroneously interpreted as data or commands. The push and pull pointers of the various clock synchronization FIFOs 102 (in the CPUs 12) and 518 (in the routers 14) need to be... ..apart, and presetting the associated FIFO queues to some known state. This done, all clock synchronization FIFOs are initialized for near frequency operation. Thus, when the system 10 is initially brought... ..in order to properly implement the lock-step operation of duplex mode operation, the clock synchronization FIFOs must be synchronized to operate with the particular source from which they receive data in order accommodate any... ..14A, 14B to the CPUs 12A, 12B must be accounted for. It is the clock synchronization FIFOs 102 of the paired CPUs 12 that operate to receive message packet symbols, adjust... ..and present symbols to the two CPUs in a simultaneous manner to maintain lock-step synchronization necessary for duplex mode operation.

In similar fashion, each symbol received by the routers 14A... ..the CPUs (which is discussed further hereinafter). Again, it is the function of the clock synchronization FIFOs 518 of the routers 14A, 14B that receive message packets from the CPUs 12... ..so that the symbols received from the two CPUs 12 are retrieved from the clock synchronization FIFOs simultaneously.

Before discussing how the clock synchronization FIFOs of the CPUs and routers are reset, initialized, and synchronized, an understanding of their operation to maintain synchronous lock-step duplex mode operation is believed helpful. Thus, referring for the moment to Fig. 23, the clock synchronization FIFOs 102 of the CPUs 12A, 12B that receive data, for example, from the router... ..underscore)Clk, from the router 14A to the CPU 12B.

Consider operation of the clock synchronization FIFOs 102x)), 102y)), to receive identical symbol streams during duplex operation. Table 6, below, illustrates... ..held by the push and pull pointer counters 128, 130 for the CPU 12A (interface unit 24a), and the content of each of the four storage locations (byte 0. ..., byte 3... ..of Table 6 show the same thing for the FIFO 102y)) of CPU 12B interface unit 24a for each symbol of the duplicated symbol stream.

Assuming the delay 640 is no... ..0" locations of the queues 126. This is because (1) the FIFOs 102 have been synchronized to operate in synchronism (a process described below), and (2) the push pointer counters 128 are clocked by the clock signal... ..of the symbol stream transmitted by the router 14A will be pulled from the clock synchronization FIFOs 102 of the CPUs 12A, 12B simultaneously, maintaining the required synchronization of received data when operating in duplex mode. In effect, the depths of the queues... ..order to achieve the operation just described with reference to Table 6, the reset and synchronization process shown in Fig 31A is used. The process not only initializes the clock synchronization FIFOs 102 of the CPUs 12A, 12B for duplex mode operation, but also operates to adjust the clock synchronization FIFOs 518 (Fig. 19A) of the CPU ports of each of the routers 14A, 14B for duplex operation. The reset and synchronization process uses the SYNC command symbol to initiate a time period, delineated by the SYNC CLK signal 970 (Fig. 31B), to reset and initialize the respective clock synchronization FIFOs of the CPUs 12A and 12B and routers 14A, 14B. (The SYNC CLK signal... ..It is of a lower frequency than that used to receive symbols by the clock synchronization FIFOs, T(underscore)Clk. For example, where T(underscore)Clk is approximately 50 MHz, the... ..signal is approximately 3.125 MHz.)

Turning now to Fig. 31A, the reset and initialization process begins at step 950 by switching the clock signals used by the CPUs 12A, 12B and routers 14A, 14B as the transmit (T(underscore)Clk) and the unit's local clock (Local Clk) clock signals so that they are derived from the same... ..In addition, configuration registers in the CPUs 12A, 12B (configuration registers 74 in the interface units 24) and the routers 14A, 14B (contained in control logic unit 509 of routers 14A, 14B) are set to the FreqLock state.

The following discussion involves step 952, and makes reference to the interface unit 24 (Fig.5), router 14A (Fig. 19A) and Figs. 31A and 31B. With the clock... ..otherwise be sent followed by a self-addressed message packet. Any message packet in the process of being received and retransmitted when the SLEEP command symbols are received and recognized by... ..per the destination address). The SLEEP command symbol operates to "quiece" router 14A for the synchronization process. The self-addressed message packet sent by the CPU 12A, when received back by the... ..message packet sent after the SLEEP command symbol would necessarily have to be the last processed by the router 14A.

At step 954 the CPU 12A checks to see if it... ..the router will assert a RESET signal 972 that is applied to the two clock synchronization FIFOs 518 contained in the input logic 5054)), 5055)) of the router that receive symbols directly from CPUs 12A, 12B. RESET, while asserted, will hold the two clock synchronization FIFOs 518 in a temporarily non-operating reset state with the push and pull pointer.....As each of the CPUs 12 receive SYNC symbols are detected by the storage and processing units of the packet receivers 96 (Figs. 5 and 6) cause the RESET signal to be asserted by the packet receivers 96 (actually, storage and processing elements 110; Fig. 6) of each CPU 12. the RESET signal is applied to the... ..t4))), CPUs 12 and routers 14A, 14B de-assert the RESET signals, and the clock synchronization FIFOs of the CPUs 12A, 12, and routers 14A, 14B are released from their reset... ..the delay,

the router 14A and CPUs 12 resume pulling data from their respective clock synchronization FIFOs and resume normal operation. The clock

synchronization FIFOs of the router 14A begin pulling symbols from the queue (previously set by RESET... ..from the CPU 12A with the T(underscore)Clk will be pushed onto the clock synchronization FIFO at, for example, queue location 0 (or whatever other location pointed to by the... ..0 (or whatever other location the push pointer was set to by RESET). The clock synchronization FIFOs of the router 14A are now synchronized to accommodate whatever delay 640 may be present in one communications path, relative to the... ..and the CPUs 12A, 12B.

Similarly, at the same virtual time, operation of the clock synchronization FIFOs 102 of both CPUs 12A, 12B is resumed, synchronizing them to the router 14A. Also, the CPUs 12A, 12B quit sending the SLEEP command... ..in favor of READY symbols, and resume message packet transmission, as appropriate.

That completes the synchronization process for the router 14A. However, the process must also be performed for the router 14B. Thus, the CPU 12A returns to step... ..however, assuming that the CPUs 12A, 12B are operating in duplex mode, the method and apparatus used to detect and handle a possible error, resulting in divergence of the CPUs from...via a message packet destined for a peripheral device of one or the other sub-processor systems 10A, 10B. Depending upon the destination of the outgoing message packet, step 1002 will... ..router 14 will issue an ERROR signal to the router control logic 509, causing the process to move to step 1004 where the router 14 detecting divergence will transmit a DVRG... ..time outs to occur. A router detecting divergence (without also detecting any simple link error) buys itself time to check the CRC of the received message packet by waiting for the... ..router 14, or received, all further message packets received from the CPUs and in the process of being routed when divergence was detected, or the DVRG symbol received, will be passed... ..1010) contained in a one of the configuration registers 74 (Fig. 5) of the interface unit 24 of each CPU.

Returning for the moment to step 1006, the determination of which... ..local" is meant to refer to the router 14A, 14B contained in the same sub-processor system 10A, 10B as the CPU. For example, referring to Fig. 1A, router 14A is... ..bit mentioned above: the bit contained in one of the configuration registers 74 of interface unit 24 (Fig. 5) of each CPU. When set to a first state, that particular CPU... ..the other CPU. In response, the state machines (not shown) within the control and status unit 509 (Fig. 19A) changes the "favorite" bits described above.

A few examples may facilitate understanding... ..DVRG symbol will echo that symbol to the routers 14A, 14B, start its internal divergence process timer, and begin determination of whether to continue or terminate. Having received a TLB symbol... ..to diverge with no errors reported. This can happen only if software (running on the processors 20) uses known divergent data to alter state. For example, suppose each CPU 12 has... ..number of the CPU 12A will differ from that of the CPU 12B. If the processors use the serial number to change the sequence of instructions executed (say, by branching if the serial number comes after some value) or to modify the value contained in a processor register, the complete "state" of the CPUs 12 will differ. In such cases, the "asymmetrical... ..of the primary CPU simply allows one CPU, and thereby the system 10, to continue processing without software intervention.

- An error at the output of the interface unit 24 of a CPU 12 will be detected by the router 14A, 14B, depending upon... ..router 14A, 14B that connects to a CPU 12 will be detected by the interface unit 24 of the affected CPU. The CPU will send a TLB symbol to the faulty... ..possible failure and, without external intervention, and transparently to the system user, remove the failing unit (CPU 12A or 12B, or router 14A or 14B) from the system to obviate or... ..reintegration." The discussion will refer to the CPUs 12A, 12B, routers 14A, 14B, and maintenance processor 18A, 18B shown forming parts of the processing system 10 illustrated in Fig. 1A. In addition, discussion will refer to the processors 20a, 20b, the interface units 24a, 24b, and the memory controllers 26a, 26b (Fig. 2) of the CPUs 12A, 12B as single units, since that is the way they function.

Reintegration is used to place two CPUs in... ..both of the paired CPUs at virtually the same time.

The major steps in the process for changing from simplex mode operation of the one on-line CPU to duplex mode... ..greater detail by the flow diagrams of Figs. 33A - 33D, generally are:

1. Setup and synchronize the two CPUs (one on-line, the other off-line) and their connected routers to... ..the memory of the on-line CPU to the off-line CPU, maintaining a tracking process that monitors changes in the memory of the on-line CPU that have not been... ..and may need to be copied over to, the off-line CPU;
3. Setup and synchronize the CPUs to run a delayed (slave) duplex mode from the same instruction stream (lock... ..will write the predetermined registers (not shown) of the control registers 74 in the interface units 24 of CPUs 12A and 12B, to a next state (after a soft operation) in... ..the off-line CPU 12B.

Next, a sequence is entered (steps 1060 - 1070) that will synchronize the clock synchronization FIFOs of the CPUs 12A, 12B and routers 14A, 14B in much the same fashion... ..the same steps described above in connection with the discussion of Figs. 31A, 31B to synchronize the clock synchronization FIFOs. The on-line CPU 12A will send the sequence of a SLEEP symbol, self-addressed message packet, and SYNC symbol which, with the SYNC CLK signal, operates to synchronize CPUs and routers. Once so synchronized, the on-line CPU 12A then, at step 1066, sends a Soft Reset (SRST) command... ..of all configuration registers and control registers (e.g., configuration registers 74 of the interface units 24) cache, and the like to memory 28 of the on-line CPU, copying the... ..time to have the system 10 off-line for reintegration. For that reason, the reintegration process is performed in a manner that allows the on-line CPU to continue executing user... ..not match that of the off-line CPU. The reason for this is that normal processing by the processor 20 of the on-line CPU can change memory content after it has been copied... ..when a memory location is written in the on-line CPU 12A during the reintegration process it is marked as "dirty;" second, all copying of memory to the off-line CPU... ..may, however, limit the ability to detect two-bit errors. But, since the memory copying process will last for a only relatively short period of time, this risk is believed acceptable...memory location in CPU 12A is made (either an incoming I/O write, or a processor write operation). The returning data (that was copied over to the off-line CPU) would... ..controller 26 (Fig. 2) of the on-line CPU to monitor memory locations in the process of being copied over to the off-line CPU 12B. The memory controller uses a... ..within the block had been written by another operation (e.g., a write by the processor 20, an I/O write, etc.), that prior write operation will flag the location in... ..still must be copied over to the off-line CPU 12B.

Returning to the reintegration process, and now to Fig. 33B, the memory tracking (AtomicWrite mechanism and using ECC to mark... ..entails writing a reintegration register (not shown; one of the configuration registers 74 of interface unit 24 - Fig. 5) to cause a reintegration (REINT) signal to be asserted. The REINT signal... ..is left alone. Throughout the incremental copy operations, the normal actions of the on-line processor will mark some memory locations dirty.

Several passes of incremental copying will need to be... ..the number of successful WriteConditional operations at the end of each pass through memory, the processors 20 can determine the effect of a given pass compared to the previous pass. When the benefits drop off, the processors 20 will give up on the precopy operations. At this point the reintegration process is ready to place the two CPUs 12A, 12B into lock-step operation.

Thus, the... ..in Fig. 33C, where at step 1100, the on-line CPU 12A momentarily halts foreground processing, i.e., execution of a user application. The remaining state (e.g., configuration registers, cache, etc.) of the on-line processors 20 and its caches is then read and written to a buffer (series of memory... ..to the off-line CPU 12B, together with a "reset vector" that will direct the processor units 20 of both CPUs 12A, 12B to a reset instruction.

Next, step 1106 will quiesce... ..to ensure that the FIFOs of the routers are clear, that the FIFOs of the processor interfaces 24 are clear, and no further incoming I/O message packets are forthcoming. At... ..symbol will be received and acted upon by both CPUs 12A, 12B, to cause the processor units 20 of each CPU to jump to the location in memory 28 containing the reset... ..a subroutine that will restore the stored state of both CPUs 12A, 12B to the

processor units 20, caches 22, registers, etc. The CPUs 12A, 12B will then begin executing the same... ..enabling of the ECC bit to mark dirty locations must now be disabled, since the processors are doing the same thing to the same memory. During this stage of the reintegration... ..encountered by CPU 12A.

Meanwhile, the bus error in the CPU 12A will cause the processor unit 20 to be forced into an error-handling routine to determine (1) the cause of... ..error was caused by an attempt to read a memory location marked dirty. Accordingly, the processor unit 20 will initiate (via the BTE 88 -- Fig. 5) the AtomicWrite mechanism to copy the... ..the SRST symbols are now received by the CPUs 12A, 12B, they will cause both processor units 20 of the CPUs to be reset to start from the same location with the... ..will periodically update, e.g., a database or audit file that is indicative of the

processing of the primary CPU up to that point in time of the update. Should the... ..in error-checking redundancy to the CPU 12B, in the same manner that the individual processor units 20a, 20b of the CPU 12A provide fail-fast, fault tolerance for the CPU - when... ..cost system is applicable, as illustrated in Fig. 34. As shown in Fig. 34, a processing system 10' includes the CPU 12A and routers 14A, 14B structured as described above. The... ..and the CPUs are also the same.

Thus, the CPU 12B' comprises only a single processor unit 20' and associated support components, including the cache 22', interface unit (IU) 24', memory controller 26', and memory 28'. Thus, while the CPU 12A is structured in the manner shown in Fig. 2, with cache processor unit, interface unit, and memory control redundancies, approximately one-half of those components are needed to implement CPU... ..stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... ..i.e., parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of... ..inputs connected to receive the data output

from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... ..for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the... ..attached, without the assistance of any other CPU 12 in the system. Many prior parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O device becomes... ..rest of the system. Other prior systems provide access to I/O through pairs of processors so that should one of the processors fail, access to the corresponding I/O is still available through the remaining I/O... ..if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... ..multiprocessing system access to every peripheral, as done here, operates to extend the "primary"/"backup" process taught in the above-identified U.S. Patent No. 4,228,496. There, a multiple CPU system may have a primary process may running on one CPU, while a backup process resides in the background on another of the CPUs. Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be... ..e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system... CPU can access any I/O without the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time... indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Specification: ...enabling of the ECC bit to mark dirty locations must now be disabled, since the processors are doing the same thing to the same memory. During this stage of the reintegration... encountered by CPU 12A.

Meanwhile, the bus error in the CPU 12A will cause the processor unit 20 to be forced into an error-handling routine to determine (1) the cause of... error was caused by an attempt to read a memory location marked dirty. Accordingly, the processor unit 20 will initiate (via the BTE 88 -- Fig. 5) the AtomicWrite mechanism to copy the... the SRST symbols are now received by the CPUs 12A, 12B, they will cause both processor units 20 of the CPUs to be reset to start from the same location with the... will periodically update, e.g., a database or audit file that is indicative of the processing of the primary CPU up to that point in time of the update. Should the... in error-checking redundancy to the CPU 12B, in the same manner that the individual processor units 20a, 20b of the CPU 12A provide fail-fast, fault tolerance for the CPU - when... cost system is applicable, as illustrated in Fig. 34. As shown in Fig. 34, a processing system 10' includes the CPU 12A and routers 14A, 14B structured as described above. The... and the CPUs are also the same.

Thus, the CPU 12B' comprises only a single processor unit 20' and associated support components, including the cache 22', interface unit (IU) 24', memory controller 26', and memory 28'. Thus, while the CPU 12A is structured in the manner shown in Fig. 2, with cache processor unit, interface unit, and memory control redundancies, approximately one-half of those components are needed to implement CPU... stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... i.e., parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of... inputs connected to receive the data output from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the... attached, without the assistance of any other CPU 12 in the system. Many prior parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O device becomes... rest of the system. Other prior systems provide access to I/O through pairs of processors so that should one of the processors fail, access to the corresponding I/O is still available through the remaining I/O... if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... multiprocessing system access to every peripheral, as done here, operates to extend the "primary"/"backup" process taught in the above-identified <PATCIT ID=PCIT0027

DNUM=US4228496A> U.S. Patent No. 4,228,496 </PATCIT>. There, a multiple CPU system may have a primary process may running on one CPU, while a backup process resides in the background on another of the CPUs. Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be... ..e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system... ..CPU can access any I/O without the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... ..four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the

processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... ..an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time... ..indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... ..When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... ..response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Claims:

1. An input/output routing apparatus for communicating digital information between a plurality of processing system elements, the communicated information being in the form of a message packet transmitted by a one of the plurality of processing system elements to another of the processing system elements, the message packet containing data identifying as the destination the another one of the plurality of processing system elements, the input/output routing apparatus comprising:

a plurality of port means each coupled to corresponding ones of the plurality of processing system elements, the port means including means bi-directionally transmitting and receiving message packets between the input/output routing apparatus and the first and second pluralities of processing system elements;

routing means responsive to the data identifying the destination for selecting a one of the plurality of port means to which the another one of the plurality of processing elements is coupled for transmitting the message packet to the another processing system element;

means including means for checking the message packet for errors, and for sending with the message packet to the another one of the plurality of processing elements indicia of the error; and

means for sending with the message packet to the another one of the plurality of processing elements indicia of an error when an error is found to exist in the message packet.

2. The input/output routing apparatus of claim 1, the routing means including means for checking the message packet for errors, and for sending with the message packet to the another one of the plurality of processing elements indicia of the lack of error in the message packet.
3. The input/output routing apparatus of claim 1, wherein the message packet contains source data identifying the one of the plurality of processing system elements as a source of the message packet, and including circuit means responsive to... ..the message packet was received at a port means coupled to the one of the processing elements.
4. The input/output routing means of claim 3, including means responsive to the... ..of error to be included in the information transmitted to the another one of the processing system elements, indicating that the message packet was not received at the port means coupled to the one of the plurality of processing system elements.
5. The input/output routing apparatus of claim 1, including means operating to receive the message packet at a one of the port means coupled to the one of the plurality of processing system elements, and for transmitting the message packet at two other of the plurality of port means simultaneously.
6. Routing apparatus for communicating digital information between a plurality of processing system elements, the digital information being in the form of message packets transmitted by a first of the plurality of processing system elements to a second of the processing system elements, the message packets including destination data identifying the second processing system element as the destination, the routing apparatus comprising:
a plurality of ports each having an input for receiving message packets and an... ..received at a predetermined input from being transmitted from a predetermined output.
7. The routing apparatus of claim 6, wherein the routing means includes addressable storage means for storing a number... ..the plurality of port means to which the message packet is communicated.
8. The routing apparatus of claim 6, the port enable means including, for each input, a register containing enable... ..means to which message packets received at such input cannot be communicated.
9. The routing apparatus of claim 8, wherein the register for each input includes a bit position corresponding to each output of the plurality of port means.
10. The routing apparatus of claim 9, wherein the bit position is set to a first digital stage to... ..of the message packet to the output corresponding to the bit position.
11. The routing apparatus of claim 7, wherein the portion of the destination data defines a region of a computing system that contains at least one of the plurality of processing system elements.
12. The routing apparatus of claim 11, wherein another portion of the destination data identifies the one processing system element within the region.
13. The routing apparatus of claim 7, including register means containing a default multi-bit entry identifying the output...

6/K/18 (Item 18 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Patent Assignee:

- Compaq Computer Corporation... ;
;;

Country	Number	Kind	Date	
Type		Pub. Date	Kind	Text
...Compaq Computer Corporation (687797) 20555 SH 249 Houston, Texas 77070-2698 US		19		
Examination...		19		
Available Text		Language	Update	Word Count
Total Word Count (Document A)				
Total Word Count (Document B)				
Total Word Count (All Documents)				

Specification: ...and 08/485,053, filed concurrently herewith.

The present invention is directed generally to data processing systems, and more particularly to a multiple processing system and a reliable system area network that provides connectivity for interprocessor and input/output... ..and communications systems to general purpose high availability commercial systems. The evolution of fault tolerant computers has been well documented (see D. P. Siewiorek, R. S. Swarz, "The Theory and Practice... ..and the Jet Propulsion laboratory began to apply fault tolerance to the development of guidance computers for aerospace applications. The 1960's also saw the development of the first AT&T electronic switching systems.

The first commercial fault tolerant machines were introduced by Tandem Computers in the 1970's for use in on-line transaction processing applications (J. Bartlett, "A NonStop Kernal," in proc. Eighth Symposium on Operating System Principles, pp... ..systems were introduced in the 1980's (O. Serlin, "Fault- Tolerant Systems in Commercial Applications," Computer, pp. 19-30, August 1984). Current commercial fault tolerant systems include distributed memory multi-processors, shared-memory transaction based systems, "pair-and-spare" hardware fault tolerant systems (see R. Freiburghouse, "Making Processing Fail-safe," Mini-micro Systems, pp. 255-264, May 1982; U.S. Patent No. 4... ..system.), and triple-modular-redundant systems such as the "Integrity" computing system manufactured by Tandem Computers Incorporated of Cupertino, California, assignee of this application and the invention disclosed herein.

Most applications of commercial fault tolerant computers fall into the category of on-line transaction processing. Financial institutions require high availability for electronic funds transfer, control of automatic teller machines, and... ..telecommunications systems.

Vendors of fault tolerant machines attempt to achieve both increased system availability, continuous processing, and correctness of data even in the presence of faults. Depending upon the particular system architecture, application software ("processes") running on the system either continue to run despite failures, or the processes are automatically restarted from a recent checkpoint when a fault is encountered. Some fault tolerant systems are provided with sufficient component redundancy to be able reconfigure around failed components, but processes running in the failed modules are lost. Vendors of commercial fault tolerant systems have extended fault tolerance beyond the processors and disks. To make large improvements in reliability, all sources of failure must be addressed... ..power supplies, fans and inter-module connections.

The "NonStop," and "Integrity" architectures manufactured by Tandem Computers Incorporated, (both respectively illustrated broadly in U.S. Patent No. 4,228,496 and U... ..assigned to the assignee of this application; NonStop and

Integrity are registered trademarks of Tandem Computers Incorporated) represent two current approaches to commercial fault tolerant computing. The NonStop system, as generally... above-identified U.S. Patent No. 4,278,496, employs an architecture that uses multiple processor systems designed to continue operation despite the failure of any single hardware component. In normal operation, each processor system uses its major components independently and concurrently, rather than as "hot backups". The NonStop system architecture may consist of up to 16 processor systems interconnected by a bus for interprocessor communication. Each processor system has its own memory which contains a copy of a message-based operating system. Each processor system controls one or more input/output (I/O) busses. Dual-porting of I/O controllers and devices provides multiple paths to each device. External storage (to the processor system), such as disk storage, may be mirrored to maintain redundant permanent data storage.

This... hardware, while fault recovery is the responsibility of the software.

Also, in the Nonstop multi-processor architecture, application software ("process") may run on the system under the operating system as "process-pairs," including a primary process and a backup process. The primary process runs on one of the multiple processors while the backup process runs on a different processor. The backup process is usually dormant, but periodically updates its state in response to checkpoint messages from the primary process. The content of a checkpoint message can take the form of complete state update, or... checkpoints were manually inserted in application programs, but currently most application code runs under transaction processing software which provides recovery through a combination of checkpoints and transaction two-phase commit protocols.

Interprocessor message traffic in the Tandem Nonstop architecture includes each processor periodically broadcasting an "I'm Alive" message for receipt by all the processors of the system, including itself, informing the other processors that the broadcasting processor is still functioning. When a processor fails, that failure will be announced and identified by the absence of the failed processor's periodic "I'm Alive" message. In response, the operating system will direct the appropriate backup processes to begin primary execution from the last checkpoint. New backup processes may be started in another processor, or the process may be run with no backup until the hardware has been repaired. U.S. Patent... example of this technique.

Each I/O controller is managed by one of the two processors to which it is attached. Management of the controller is periodically switched between the processors. If the managing processor fails, ownership of the controller is automatically switched to the other processor. If the controller fails, access to the data is maintained through another controller.

In addition to providing hardware fault tolerance, the processor pairs of the above-described architecture provide some measure of software fault tolerance. When a processor fails due to a software error, the backup processor frequently is able to successfully continue processing without encountering the same error. The software environment in the backup processor typically has different queue lengths, table sizes, and process mixes. Since most of the software bugs escaping the software quality assurance tests involve infrequent data dependent boundary conditions, the backup processes often succeed.

In contrast to the above-described architecture, the Integrity system illustrates another approach... fault recovery is the logical choice since few modifications to the software are required. The processors and local memories are configured using triple-modular-redundancy (TMR). All processors run the same code stream, but clocking of each module is independent to provide tolerance... three streams is asynchronous, and may drift several clock periods apart. The streams are re-synchronized periodically and during access of global memory. Voters on the TMR Controller boards detect and mask failures in a processor module. Memory is partitioned between the local memory on the triplicated processor boards and the global memory on the duplicated TMRC boards. The duplicated portions of the... techniques to detect failures. Each global memory is dual ported and is interfaced to the processors as well to the I/O Processors (IOPs). Standard VME peripheral controllers are interfaced to a pair of busses through a Bus...

...the BIMs to switch control of all controllers to the remaining IOP. Mirrored disk storage units may be attached to two different VME controllers.

In the Integrity system all hardware failures... ..reintegrated on-line.

The preceding examples illustrate present approaches to incorporating fault tolerance into data processing systems. Approaches involving software recovery require less redundant hardware, and offer the potential for some... ..have been developed on other systems.

Thus, the systems described above provide fault tolerant data processing either by hardware (e.g, fail-functional, employing redundancy) or by software techniques (fail-fast... ..hardware). However, none of the systems described are believed capable of providing fault tolerant data processing, using both hardware (fail-functional) and software (fail-fast) approaches, by a single data processing system.

Computing systems, such as those described above, are often used for electronic commerce: electronic data interchange (EDI) and global messaging. Today's demands upon such electronic commerce, however, is demanding more and more throughput capacity as the number of users increases and... ..networks such as local area networks (LANS), and the like.

A key requirement for a server architecture is the ability to move massive quantities of data. The server should have high bandwidth that is scalable, so that added throughput capacity can be added... ..response time, latency affects service levels and employee productivity.

The present invention provides a multiple-processor system that combines both of the two above-described approaches to fault tolerant architecture, hardware redundancy and software recovery techniques, in a single system.

Broadly, the present invention includes a processing system composed of multiple sub-processing systems. Each sub-processing system has, as the main processing element, a central processing unit (CPU) that in turn comprises a pair of processors operating in lock-step, synchronized fashion to execute each instruction of an instruction stream at the same time. Each of the sub-processing systems further include an input/output (I/O) system area network system that provides redundant communication paths between various components of the larger

processing system, including a CPU and assorted peripheral devices (e.g., mass storage units, printers, and the like) of a sub-processing system, as well as between the sub-processors that may make up the larger overall processing system. Communication between any component of the processing system (e.g., a CPU and another CPU, or a CPU and any peripheral device, regardless of which sub-processing system it may belong to) is implemented by forming and transmitting packetized messages that are... ..responsible for choosing the proper or available communication paths from a transmitting component of the processing system to a destination component based upon information contained in the message packet. Thus, the... ..peripherals, but permits it to also be used for interprocessor communications.

As indicated above, the processing system of the present invention is structured to provide fault-tolerant operation through both "fail... ..at a variety of points in the various data paths between the (lock-step operated) processor elements of the CPU and its associated memory. In particular, the processing system of the present invention conducts error-checking at an interface, and in a manner... ..little impact on performance. Prior art systems typically implement error-checking by running pairs of processors, and checking (comparing) the data and instruction flow between the processors and a cache memory. This technique of error-checking tended to add delay to the... ..error-checking precluded use of off-the-shelf parts that may be available (i.e., processor/cache memory combinations on a single semiconductor chip or module). The present invention performs error-checking of the processors at points that operate at slower rates, such as the main memory and I/O interfaces which operate at slower speeds than the processor-cache interface. In addition, the error-checking is performed at locations that allow detection of errors that

may occur in the processors, their cache memory, and the I/O and memory interfaces. This allows simpler designs for... other data integrity checks.

Error-checking of the communication flow between the components of the processing system is achieved by adding a cyclic-redundancy-check (CRC) to the message packets that... "Good" (TPG) or "This Packet Bad" (TPB) - is appended to every packet. A maintenance diagnostic processor can use this information to isolate a link or router element that introduces an error... of topologies, so that alternate paths can be provided between any two elements of a processing system (e.g., between a CPU and an I/O device), for communication in the... so (e.g., by creating a "deadlock" condition, discussed further below).

The CPUs of a processing system are capable of operating in one of two basic modes: a "simplex mode" in... independently of the other, or a "duplex" mode in which pairs of CPUs operate in synchronized, lock-step fashion. Simplex mode operation provides the capability of recovering from faults that are... U.S. Pat. No. 4,228,496 which teaches a multiprocessing system in which each processor has the capability of checking on the operability of its sibling processors, and of taking over the processing of a processor found or believed to have failed). When operating in duplex mode, the paired CPUs both... fault tolerant platform for less robust operating systems (e.g., the UNIX operating system). The processing system of the present invention, with the paired, lock-step CPUs, is structured so that a fault), primarily through hardware.

When the processing system is operating in duplex mode, each CPU pair uses the I/O system to access any peripheral of the processing system, regardless of which (of the two, or more) sub-processor system the peripheral may be ostensibly a member of. Also, in duplex mode, message packets... message for the CPU pair (from either a peripheral device such as a mass storage unit or from a processing unit), will replicate the message and deliver it to both CPUs of the pair using synchronization methods that ensure that the CPUs remain synchronized. In effect, the duplex CPU pair, as viewed from the I/O system and other... as a single CPU. Thus, the I/O system, which includes elements from all sub-processing systems, is made to be seen by the duplex CPU pair as one homogeneous system... a multiprocessor system in which the CPU of any one is actually a pair of synchronized, lock-step CPUs.

Yet another important aspect of the present invention is that interrupts issuing... interrupts via the message packet system ensures that they will arrive at duplexed CPUs in synchronized fashion, in the same manner as I/O message packets. Interrupt message packets will contain... the system. In addition, using the same messaging system to communicate data between I/O units and the CPUs and to communicate interrupts to the CPUs preserves the ordering of I... the implementation of a technique of validating access to the memory of any CPU. The processing system, as structured according to the present invention, permits the memory of any CPU to... to handle input/output information transfers between a CPU and any other component of the processor system. Thereby, the individual processor units of the CPU are removed from the more mundane tasks of getting information from memory and out onto the TNet network, or accepting information from the network. The processor unit of the CPU merely sets up data structures... is required, where in memory the response is to be placed when received. When the processor unit completes the task of creating the data structure, the block transfer engine is notified to... response is received, it is routed to the expected memory location identified, and notifies the processor unit that the response was received.

Further aspects and features of the present invention will become... invention, which should be taken in conjunction with the accompanying drawings.

Fig. 1A illustrates a processing system constructed in accordance with the teachings of the present invention, and Figs. 1B and 1C illustrate two alternate configurations of the processing system of Fig. 1A, employing clusters or arrangements of the processing system of Fig. 1A;

Fig. 2 illustrates, in simplified block diagram form, the central processing unit (CPU) that forms a part of each sub-processor system of Figs. 1A - 1C;

Figs. 3A - 3D and 4A - 4C illustrate the construction of... ..the area network I/O system shown in Fig. 2;

Fig. 5 illustrates the interface unit that forms a part of the CPUs of Fig. 2 to interface the processor and memory with the I/O area network system;

Fig. 6 is a block diagram, illustrating a portion of packet receiver of the interface unit of Fig. 5;

Fig. 7A diagrammatically illustrates the clock synchronization FIFO (CS FIFO) used by the packet receiver section packet receiver shown in Fig. 6;

Fig. 7B is a block diagram of a construction of the clock synchronization FIFO structure shown in Fig. 7A;

Fig. 8 illustrates the cross-connections for error-checking outbound transmissions from the two interface units of a CPU;

Fig. 9 illustrates an encoded (8B to 9B) data/command symbol;

Fig. 10 illustrates the method and structure used by the interface unit of Fig. 5 to cross-check for errors data being transferred to the memory controllers... ..of a CPU of Fig. 2 to other (external to the CPU) components of the processing system;

Fig. 12 is a block diagram that diagrammatically illustrates the formation of an address... ..14A illustrates the logic for posting interrupt requests to queues in memory and to the processor units of the CPU of Fig. 2;

Fig. 14B illustrates the process used to form a memory address for a queue entry;

Fig. 15 is a block... ..data output constructs formed in the memory of the CPU of Fig. 2 by a processor unit, and containing data to be sent via the area I/O networks shown in Figs. 1A - 1C, and also illustrating the block transfer engine (BTE) unit of the interface unit of Fig. 5 that operates to access the data output constructs for transmission to the... ..pair of memory controllers between memory of a CPU of Fig. 2 and its interface unit for accessing from memory 72 bits of data, including two simultaneously-accessed 32-bit words... ..other for error-checking;

Fig. 19A is a simplified block diagram illustration of the router unit used in the area input/output networks of the processing systems shown in Figs. 1A - 1C;

Fig. 19B illustrates comparison on two port inputs of the router unit of Fig. 19A;

Fig. 20A is a block diagram the construction of one of the six input ports of the router unit shown in Fig. 19A;

Fig. 20B is a block diagram of the synchronization logic used to validate command/data symbols received at an input port of the router unit of Fig. 19A;

Fig. 21A is a block diagram illustration of the target port selection... ..is a block diagram illustration of one of the six output ports of the router unit shown in Fig. 19A;

Fig. 23 is an illustration of the method used to transmit identical information to a duplexed pair CPUs of Fig. 2 in synchronized fashion when the processing system is operating in lock-step (duplex) mode, using a pair the FIFOs of Fig... ..is a simplified block diagram illustrating the clock generation system of each of the sub-processing systems of Figs. 1A - 1C for developing the plurality of clock signals used to operate the various elements of that sub-processing system;

Fig. 25 illustrates the topology used to interconnect the clock generation systems of paired sub-processing systems for synchronizing the various clock signals of the pair of sub-processing systems to one another;

Fig. 26A and 26B illustrates a FIFO constant rate clock control logic used to control the clock synchronization FIFO of Figs. 8 or 20 in the situation when the two clocks used to... ..structure of the on-line access port (OLAP) used to provide access to the maintenance

processor (MP) to the various elements of the system of Fig. 1A (or those of Figs... ..the soft-flag logic used to handle asymmetric variables between the CPUs of paired sub-processor systems operating in duplex mode;

Fig. 31A shows a flow diagram, and Fig. 31B illustrates a portion of SYNC CLK, both of which are used to reset and synchronize the clock synchronization FIFOs of the CPUs and routers of the processing system of Fig. 1A that receive information from each other;

Fig. 32 is a flow... ..33A - 33D generally illustrate the procedure used to bring an one of the CPUs of processing system shown in Fig. 1A into lock-step, duplex mode operation with the other of the CPUs without measurably halting operation of the processing system; and

Fig. 34 illustrates a reduced cost architecture incorporating teachings of the invention; and... ..to the figures and, for the moment, principally Fig. 1A, there is illustrated a data processing system, designated with the reference 10, constructed according to the various teachings of the present invention. As Fig. 1A shows, the data processing system 10 comprises two sub-processor systems 10A and 10B each of which are substantially the same in structure and function... ..should be appreciated that, unless noted otherwise, a description of any one of the sub-processor systems 10 will apply equally to any other sub-processor system 10.

Continuing with Fig. 1A therefore, each of the sub-processor systems 10A, 10B is illustrated as including a central processing unit (CPU) 12, a router 14, and a plurality of input/output (I/O) packet interfaces... ..one of the I/O packet interfaces 16 will also have coupled thereto a maintenance processor (MP) 18.

The MP 18 of each sub-processor system 10A, 10B connects to each of the elements of that sub-processor system via an IEEE 1149.1 test bus 17 (shown in phantom in Fig. 1A... ..accompanying clock signal. As Fig. 1A further illustrates, TNet Links L also interconnect the sub-processor systems 10A and 10B to one another, providing each sub-processor system 10 with access to the I/O devices of the other as well as inter-CPU communication. As will be seen, any CPU 12 of the processing system 10 can be given access to the memory of any other CPU 12, although... ..the memory of a CPU 12 by a wayward peripheral device 17.

Preferably, the sub-processor systems 10A/10B are paired as illustrated in Fig. 1A (and Figs 1B and 1C, discussed below), and each sub-processor system 10A/10B pair (i.e., comprising a CPU 12, at least one router 14... ..12A) connects, by a TNet Link L to a router (14A) of the corresponding sub-processor system (e.g., 10A). Conversely, the Y port connects the CPU (12A) to the router (14B) of the companion sub-processor system (10B). This latter connection not only provides a communication path for access by a CPU (12A) to the I/O devices of the other sub-processor system (10B), but also to the CPU (12B) of that system for inter-CPU communication.

Information is communicated between any element of the processing system 10 and any other element (e.g., CPU 12A of sub-processor system 10A) of the system and any other element of the system (e.g., an I/O device associated with an I/O packet interface 16B of sub-processor system 10B) via message "packets." Each message packet is made up of a number of... ..this reason, a unique method of receiving the symbols at the receiver, using a clock synchronization first-in-first-out (CS FIFO) storage structure (described more fully below), has been developed... ..locked operation means just that: the frequencies of the clock signals of the transmitter and receiver units are locked, although not necessarily in phase. Frequency locked clock signals are used to transmit symbols between the routers 14A, 14B and the CPUs 12 of paired sub-processor systems (e.g., sub-processor systems 10A, 10B, Fig. 1A). Since the clocks of the transmitting and receiving element are not phase related, a clock synchronization FIFO is again used -- albeit operating in a slightly different mode from that used for... ..difference, as will be seen, is due to the fact that pairs of the sub-processor systems 10 can be operated in a synchronized, lock-step mode, called duplex mode, in which each CPU 12 operates to execute the... ..1A illustrates another feature of the invention: a cross-link connection between the two sub-processor systems 10A, 10B through the use of additional routers 14 (identified in Fig. 1A as... ..added routers RX1)), RX2)), RY1)), and RY2)) form a cross-link connection between

the sub-processors 10A, 10B (or, as shown, "sides" X and Y, respectively) to couple them to I... ..shown in Fig. 1A, the routers RX2)) and RY2)) provide the I/O packet interface units 16x and 16y with a dual ported interface. Of course, it will now be evident... ..lend themselves to being used in a manner that can extend the configuration of the processing system 10 to include additional sub-processor systems such as illustrated in Figs. 1B and 1C. In Fig. 1B, for example, one... ..of each of the routers 14A and 14B is used to connect the corresponding sub-processor systems 10A and 10B to additional sub-processor systems 10A' and 10B' forming thereby a larger processing system comprising clusters of the basic processing system 10 of Fig. 1.

Similarly, in Fig. 1C the above concept is extended to form an eight sub-processor system cluster, comprising sub-processor systems pairs 10A/10B, 10A'/10B', 10A"/10B", and 10A'''/10B'''. In turn, each of the sub-processor systems (e.g., sub-processor system 10A) will have essentially the same basic minimum configuration of a CPU 12, a... ..by a I/O packet interface 16, except that, as Fig. 1C shows, the sub-processor systems 10A and 10B include additional routers 14C and 14D, respectively, in order to extend the cluster beyond sub-processor systems 10A'/10B' to the sub-processor systems 10A"/10B" and 10A'''/10B'''. As Fig. 1C further illustrates, unused ports 4 and... ..the routers 14 when configuring the topology of the system 10, any CPU 12 of processing system 10 of Fig. 1C can access any other "end unit" (e.g., a CPU or I/O device) of any of the other sub-processor systems. Two paths are available from any CPU 12 to the last router 14 connecting to the I/O packet interface 16. For example, the CPU 12B of the sub-processor system 10B' can access the I/O 16''' of sub-processor system 10A''' via router 14B (of sub-processor system 10B'), router 14D, and router 14B (of sub-system 10B''') and, via link LA... ..10A'''), OR via router 14A (of sub-system 10A'), router 14C, and router 14A (sub-processor system 10A'''). Similarly, CPU 12A of sub-processor system 10A" may access (via two paths) memory contained in the CPU 12B of sub-processor 10B to read or write data. (Memory accesses by one CPU 12 of another component of the processing system requires, as will be seen, the components seeking access to have authorization to do... ..prevents corruption of memory data of a CPU by erroneous access.)

The topology of the processing system shown in Fig. 1B is achieved by using port 1 of the routers 14A, 14B, and auxiliary TNet links LA, to connect to the routers 14A', 14B' of sub-processor systems 10A', 10B'. The topology thereby obtained establishes redundant communication paths between any CPU 12 (12A, 12B, 12A', 12B') and any I/O packet interface 16 of the processing system 10 shown in Fig. 1B. For example, the CPU 12A' of the sub-processor system 10A' may access the I/O 16A of sub-processor system 10A by a first path formed by the router 14A' (in port 4, out... ..shown in Fig. 1B. By interconnecting one port of each router 14 of each sub-processor pair, and using additional auxiliary TNet links LA (illustrated in Fig. 1C with the dotted line connections) between the ports 1 of the routers 14 (14A" and 14B") of sub-processor systems 10A", 10B" and 10A''', 10B''', two separate, independent data paths can be found between any CPU 12 and any I/O packet interface 16. In this fashion, any end unit (i.e., a CPU 12 or an I/O packet interface 16) will have at least two paths to any other end unit.

Providing alternate paths of access between any two end units (e.g., between a CPU 12 and any other CPU 12, or between any CPU... ..any two of the remaining fault domains. Here, a fault domain could be a sub-processor system (e.g., 10A). Thus, if the sub-processor system 10A were brought down because of a failure the electrical power being supplied, without... ..TNet link LA between the routers 14A''' and 14B''', the CPU 12B of the sub-processor system 10B would have lost access to the I/O packet interface 16''' (via router... ..with the loss of the router 14A (and router 14C) by loss of the sub-processor system 10A, communications between the CPU 12B is still possible via the route of router... ..equally to CPU 12B. As Fig. 2 shows, the CPU 12A includes a pair of processor units 20a, 20b that are configured for synchronized, lock-step operation in that both processor units 20a, 20b receive and execute identical instructions, and issue identical data and command outputs, at substantially the same moments in time. Each of the processor units 20a and 20b is connected, by a bus 21 (21a, 21b) to a corresponding cache memory 22. The particular type of processor units used could contain sufficient internal cache memory so that the cache memory 22 would not... ..22 could be used to supplement any cache memory that may be internal to the processor units 20. In any event, if the cache memory 22 is used, the bus 21 is...22 address bits, 3 bits of parity covering the address, and 7 control bits.

The processors 20a, 20b are also respectively coupled, via a separate 64-bit address/data bus 23 to X and Y interface units 24a, 24b. If desired, the address/data communicated on each bus 23a, 23b could also be protected by parity, although this will increase the width of the bus. (Preferably, the processors

20 are constructed to include RISC R4000 type microprocessors, such as are available from the MIPS Division of Silicon Graphics, Inc. of Santa Clara, California.)

The X and Y interface units 24a, 24b operate to communicate data and command signals between the processor units 20a, 20b and a memory system of the CPU 12A, comprising a memory controller (MC... ..MC halves 26a and 26b) and a dynamic random access memory array 28. The interface units 24 interconnect to each other and to the Mcs 26a, 26b by a 72-bit... ..accompanied by 8 bits of ECC) are written to the memory 28 by the interface units 24, one interface unit 24 will drive only one word (e.g., the 32 most significant portion) of the doubleword being written while the other interface unit 24 writes the other word of the double word (e.g., the least significant 32-bit portion of the doubleword). In addition, on each write operation the interface units 24a, 24b perform a cross-check operation on the data not written by that interface unit 24 with the data written by the other to check for errors; on read operations... ..accessed corresponds to the address of the location from which the doubleword was stored.

Interface units 24a, 24b of the CPU 12A form the circuitry to respectively service the X and Y (I/O) ports of the CPU 12A. Thus, the X interface unit 24a connects by the bi-directional TNet Link Lx to a port of the router 14A of the processor system 10A (Fig. 1A) while the Y interface unit 24b similarly connects to the router 14B of the processor system 10B by TNet Link Ly. The X interface unit 24a handles all I/O traffic between the router 14A and the CPU 12A of the sub-processor system 10A. Likewise, the Y interface unit 24b is responsible for all I/O traffic between the CPU 12A and the router 14B of companion sub-processor system 10B.

The TNet Link Lx connecting the X interface unit 24a to the router 14A (Fig. 1) comprises, as above indicated, two 10-bit buses... ..bus 32x)) carries data incoming from the router 14A. In similar fashion, the Y interface unit 24b is connected to the router 14B (of the sub-processor system 10B) by two 10-bit busses: 30y)) (for outgoing transmissions) and 32y)) (for incoming transmissions), together forming the TNet Link Ly.

The X and Y interface units 24a, 24b are synchronously operated in lock-step, performing substantially the same operations at substantially the same times. Thus, although only the X interface unit 24a actually transmits data onto the bus 30x)), the same output data is being produced by the Y interface unit 24b, and used for error-checking. The Y interface unit 24b output data is coupled to the X interface unit 24a by a cross-link 34y)) where it is received by the X interface unit 24a and compared against the same output data produced by the X interface unit. In this way the outgoing data made available at the X port of the CPU... ..the port of the CPU 12A is checked. The output data from the Y interface unit 24b is coupled to the Y port by a 10-bit bus 30y)), and also to the X interface unit 24a by the 9-bit cross-link 34y)) where is checked with that produced by the X interface unit.

As mentioned, the two interface units 24a, 24b operate in synchronous, lock-step with one another, each performing substantially the same... ..X and/or Y ports of the CPU 12A must be received by both interface units 24a, 24b to maintain the two interface units in this lock-step mode. Thus, data received by one interface unit 24a, 24b is passed to the other, as indicated by the dotted lines and 9... ..connections 36x)) (communicating incoming data being received at the X port by the X interface unit 24a to the Y interface unit 24b) and 36y)) (communicating data received at the Y port by the Y interface unit 24b to the X interface unit 24a).

Certain more robust operating systems are structured with a fault-tolerant capability in the... ..example, U.S. Patent No. 4,817,091 teaches a multiprocessor system in which each processor periodically messages each of the processors of the system (including itself), under software control, to thereby provide an indication of continuing operation. Each of the processors, in addition to performing its normal tasks, operates as a backup processor to another of the processors. In the event one of the backup processors fails to receive the messaged indication from a

sibling processor, it will take over the operation of that sibling (now thought to be inoperative), in... platform for both types of software. Thus, when a robust operating system is available, the processing system 10 can be configured to operate in a "simplex" mode in which each of... left, in most instances, to software.

Alternatively, for less robust operating systems and software, the processing system 10 provides a hardware-based fault-tolerance by being configured to operate in a... g., CPUs 12A, 12B) are coupled together as shown in Fig. 1A, to operate in synchronized, lock-step fashion, executing the same instructions at the substantially the same moment in time... data and command symbols. In order to simplify the design of the CPU 12, the processors 20 are precluded from communicating directly with any outside entity (e.g., another CPU 12... O device via the I/O packet interface 16). Rather, as will be seen, the processor will construct a data structure in memory and turn over control to the interface units 24. Each interface unit 24 includes a block transfer engine (BTE; Fig. 5) configured to provide a form of... to the destination according to information contained in the message packet.

The design of the processing system 10 permits a memory 28 of a CPU to be read or written by... via the routers 14. Accordingly, before continuing with the description of the construction of the processing system 10, it would be of advantage to understand first the configuration of the data... information.

As indicated, the HADC message packet operates to communicate write data between the end units (e.g., CPU 12) of the processing system 10. Other message packets, however, may be differently constructed because of their function and... CRC. The HC message packet is used to acknowledge a request to write data.

Interface Unit:

The X and Y interface units 24 (i.e., 24a and 24b - Fig. 2) operate to perform three major functions within the CPU 12: to interface the processors 20 to the memory 28; to provide an I/O service that operates transparently to, but under the control of, the processors; and to validate requests for access to the memory 28 from outside sources.

Regarding first the interface function, the X and Y interface units 24a, 24b operate to respectively communicate processors 20a, 20b to the memory controllers (Mcs 26a, 26b) and memory 28 for writing and... fast checking of the data read/written. For example, write operations have the two interface units 24a, 24b cooperating to cross-check the data to be written to ensure its integrity (and at the same time, the interface units 24 will operate) to develop an error correcting code (ECC) that covers, as will be... to have been retrieved from the appropriate address.

With respect to I/O access, the processors 20 are not provided with the ability to communicate directly with the input/output systems... must write data structures to the memory 28 and then pass control to the interface units 24 which perform a direct memory access (DMA) operation to retrieve those data structures, and... indicated in the data structure itself.)

The third function of the X and Y interface units 24, access validation to the memory 28, uses an address validation and translation (AVT) table maintained by the interface units. The AVT table contains an address for each system component (...the incoming message packets are virtual addresses. These virtual addresses are translated by the interface unit to physical addresses recognizable by the memory control units 26 for accessing the memory 28.

Referring to Fig. 5, illustrated is a simplified block diagram of the X interface unit 24a of the CPU 12A. The companion Y interface unit 24b (as well as the interface units 24 of the CPU 12B, or any other CPU 12) is of substantially identical construction. Accordingly, it will be understood that a description of the interface unit 24a will apply equally to the other interface units 24 of the processing system 10.

As Fig. 5 illustrates, the X interface unit 24a includes a processor interface 60, a memory interface 70, interrupt logic 86, a block transfer engine (BTE) 88, access validation and translation logic 90, a packet transmitter 94, and a packet receiver 96.

Processor Interface:

The processor interface 60 handles the information flow (data and commands) between the processor 20a and the X interface unit 24a. A processor bus 23, including a 64 bit address and data bus (SysAD) 23a and a 9 bit command bus 23b, couples the processor 20a and the processor interface 60 to one another. While the SysAD bus 23a carries memory address and data... ..and qualifying commands carried at substantially the same time on the SysAD bus 23a. The processor interface 60 operates to interpret commands issued by the processor unit 20a in order to pass reads/writes to memory or control registers of the processor interface. In addition, the processor interface 60 contains temporary storage (not shown) for buffering addresses and data for access to... ..26). Data and command information read from memory is similarly buffered en route to the processor unit 20a, and made available when the processor unit is ready to accept it. Further, the processor interface 60 will operate to generate the necessary interrupt signalling for the X interface unit 24a.

The processor interface 60 is connected to a memory interface 70 and to configuration registers 74 by a bi-directional 64 bit processor address/data bus 76. The configuration registers 74 are a symbolic representation of the various control registers contained in other components of the X interface

unit 24a, and will be discussed when those particular components are discussed. However, although not specifically... ..throughout other of the logic that is used to implement the X interface 24a, the processor address/data bus 76 is likewise coupled to read or write to those registers.

Configuration registers 74 are read/write accessible to the processor 20a; they allow the X interface unit to be "personalized." For example, one register identifies the node address of the CPU 12A... ..with the CPU 12A; another, readable only, contains a fixed identification number of the interface unit 24, and still other registers define areas of memory that can be used by, for... ..logic 90, etc.) employing them are discussed.

The memory interface 70 couples the X interface unit 24a to the memory controllers 26 (and to the Y interface unit 24b; see fig. 2) by a bus 25 that includes two 36 bi-directional bit... ..25a, 25b. The memory interface operates to arbitrate between requests for memory access from the processor unit 20, the BTE 88, and the AVT logic 90. In addition to memory accesses from the processor unit 20a, the memory 28 may also be accessed by components of the processing system 10 to, for example, store data requested to be read by the processor unit 20a from an I/O unit 17, or memory 28 may also be accessed for I/O data structures previously set up in memory by the processor unit. Since these accesses are all asynchronous, they must be arbitrated, and the memory interface 70... ..command information accessed from the memory 28 is coupled from the memory interface to the processor interface 60 by a memory read bus 82, as well as to an interrupt logic... ..doubleword quantities. However, while the memory interfaces 70 of both the X and Y interface units 24a and 24b formulate and apply the (64-bit) doubleword to the bus 25, each... ..by the memory interface 70 are coupled to the memory interface by the companion interface unit 24 where they are compared with the same 32 bits for error.

Digressing for the... ..containing interrupt information are received, that information is conveyed to the interrupt logic 86 for processing and posting for action by the processor 20, along with any interrupts generated internal to the CPU 12A. Internally generated interrupts will... ..register 71 (internal to the interrupt logic 86), indicating the cause of the interrupt. The processor 20 can then read and act upon the interrupt. The interrupt logic is discussed more fully below.

The BTE 88 of the X interface unit 24a operates to perform direct memory accesses, and provides the mechanism that allows the processors 20 to access external resources. The BTE 88 can be set-up by the processors 20 to generate I/O requests, transparent to the processors 20 and notify the processors when the requests are complete. The BTE logic 88 is discussed further below.

Requests for... ..8 byte wide format necessary for storing in the memory 28.

Outgoing message packets containing processor originated transaction requests (e.g., a read request asking for a block data from an I/O unit) are monitored by the request transaction logic (RTL) 100. The RTL 100 provides a time... ..will generate an interrupt (handled and reported by the interrupt logic 86) to inform the processor 20 that the request was not honored. In addition, the RTL 100 will validate responses... ..28 (by the DMA operation of the BTE 86) at a location known to the processor 20 so that it can locate the response.

Each of the CPUs 12 are checked... ..discussed. One such check is an on-going monitor of the operation of the interface units 24a, 24b of each CPU. Since the interface units 24a, 24b operate in lock-step synchronism checking can be performed by monitoring the operating states of the paired interface units 24a, 24b by a continuous comparison of certain of their internal states. This approach is implemented by using one stage of a state machine (not shown) contained in the unit 24a of CPU 12A, and comparing each state assumed by that stage with its identical state machine stage in the interface unit 24b. All units of the interface units 24 use state machines to control their operations. Preferably, therefore, a state machine of the memory interface 70 that controls the data transfers between the interface unit 24 and the MC 26 is used. Thus, a selected stage of the state machine used in the memory interface 70 of the interface unit 24a is selected. An identical stage of a state machine of one of the interface unit 24b is also selected. The two selected stages are communicated between the interface units 24a, 24b and received by a compare circuit contained in both interface units 24a, 24b. As the interface units operate lock-step with one another, the state machines will likewise march through the same identical states, assuming each state at substantially the same moments in time. If an interface unit encounters an error, or fails, that activity will cause the interface units to diverge, and the state machines will assume different states. The time will come when... ..that will bring to the attention of the CPUs 12A (or 12B) that the interface units 24a, 24b of that CPU are no longer in lock-step, and to act accordingly... ..X port, receiving only those message packets transmitted by the router 14A of the sub-processor system 10A (Fig. 1A). The Y port is serviced by the Y interface unit 24b to receive message packets from the router 14B of the companion sub-processor system 10B. However, both interfaces (as well as Mcs 26 and processor 20), as has been indicated, are basically mirror images of one another in that both... ..in both structure and function. For this reason, message packet information, received by one interface unit (e.g., 24a) must be passed for processing also to the companion interface unit (e.g., 24b). Further, since both interface units 24a, 24b will assemble the same message packets for transmission from the X or the Y ports, the message packet being transmitted by the interface unit (e.g., 24b) actually being communicated ...associated port (e.g., the Y port) will also be coupled to the other interface unit (e.g., 24a) for cross-checking for errors. These features are illustrated in Figs. 6... ..receiving portions of the packet receivers 96 (96x, 96y) of the X and Y interface units 24a, 24b are broadly illustrated. As shown, each packet receiver 96x, 96y has a clock... ..receive a corresponding one of the TNet Links 32. The CS FIFOs 102 operate to synchronize the incoming command/data symbols to the local clock of the packet receiver 96, buffering... ..104x, coupled to the MUX 104y of the packet receiver 96y of the Y interface unit 24b by the cross-link connection 36x)). In similar fashion, information received at the Y port is coupled to the X interface unit 24a by the cross-link connection 36y)). In this manner, the command/data symbols of... ..packets received at one of the X, Y ports by the corresponding X, Y, interface unit 24a, 24b is passed to the other so that both will process and communicate the same information on to other components of the interface units 24 and/or memory 28.

Continuing with Fig. 6, depending upon which port X, Y... ..or the other of the CS FIFOs 102x, 102y for communication to the storage and processing logic 110 of the interface unit 24. The information contained in each 9-bit symbol is an 8-bit byte of... ..the encoding of which is discussed below with respect to Fig. 9. The storage and processing logic 110 will first translate the 9-bit symbols to 8-bit data or command... ..the outputs of the CS FIFOs 102x, 102y are also coupled to a command decode unit in addition to the MUX 104. The command decode unit operates to recognize command symbols (differentiating them from data symbols in a manner that is... ..below), decoding them to generate therefrom command signals that are applied to a receiver control unit, a state machine-based element that functions to control packet receiver operations.

As indicated above... ..at the output of the MUX 104, the receiver control portion of the storage control unit enables CRC check logic 106 to calculate a CRC symbol while the data symbols are... ..below, CS FIFOs are found not only

in the packet receivers 96 of the interface units 24, but also at each receiving port of the routers 14 and the I/O... ..an even more important part, and perform a unique function, when a pair of sub-processor systems are operating in duplex mode and the two CPUs 12A and 12B of the sub-processor systems 10A, 10B operate in synchronized, lock-step, executing the same instructions at the same time. When operating in this latter...difficult to ensure that the clocking regime of the routers 14A and 14B are exactly synchronized to those of the CPUs 12A and 12B - even when using frequency locked clocking. In... ..used to transmit symbols to a CPU 12 and the clock used by an interface unit 24 to receive those symbols.

The structure of the CS FIFO 102 is diagrammatically illustrated... ..i.e., a packet) or IDLE symbols - except during certain situations (e.g., reset, initialization, synchronization and others discussed below). As explained above, each symbol held in the transmit register 120... ..same symbol leaving the storage queue, allowing each symbol entering the storage queue 126 to settle before it is clocked out and passed to the storage and processing units 110x (and 110y) by the MUX 104x (and 104y). Since the transmit and receive clocks...functioning in duplex mode) operate to transmit symbols with near frequency clocking. Even so, clock synchronization FIFOs are used at these other ports to receive symbols transmitted with near frequency clocking, and the structure of these clock synchronization FIFOs are substantially the same as that used in frequency locked environments, i.e., that... ..of the storage queue 126 are nine bits wide; in near frequency environments, the clock

synchronization FIFOs use symbol locations of the queue 126 that are 10 bits wide, the extra... ..the faster clock source. To handle this clock drift, the two pointers are effectively re-synchronized periodically.

When the CPUs 12 are paired and operating in duplex mode, all four interface units 24 operate in lock-step to, among other things, transmit the same data and receive... ..simplex mode, each independent of the other, clocking need only be near frequency.

The interface unit 24 receives a SYNC CLK signal that is used in combination with a SYNC command symbol to initialize and synchronize the Rcv register 124 to the transmitting router 14. When using either near frequency or... ..102X preferably begin from some known state. Incoming symbols are examined by the storage and processing units 110 of the packet receivers 96. The storage and processing units look for, and act upon as appropriate, command symbols. Pertinent here is that when the... ..receives a SYNC command symbol it will be decoded and detected by the storage and processing unit 110. Detection of the SYNC command symbol by the storage and processing unit 110 causes assertion of a RESET signal. The RESET signal, under synchronous control of the SYNC CLK signal, is used to reset the input buffers (including the clock synchronization buffers) to predetermined states, and synchronize them to the routers 14.

The synchronization of the CS FIFOs 102 of the interface units 24 those of one or both routers 14A, 14B is discussed more fully below in the section discussing synchronization.

Packet Transmitter:

Each interface unit 24 is assigned to transmit from and receive at only one of the X or Y ports of the CPU 12. When one of the interface units 24 transmits, the other operates to check the data being transmitted. This is an important... ..shows, in abbreviated form, the packet transmitters 94x, 94y of the X and Y interface units 24a, 24b, respectively. Both packet transmitters are identically constructed, so that discussion of one (packet ...logic 152 that receives, from the BTE 88 or AVT 90 of the associated interface unit (here, the X interface unit 24a) the data to be transmitted - in doubleword (64-bit) format. The packet assembly logic... ..and Y ports: they are either symbols that make up a message packet in the process of being transmitted, or IDLE symbols, or other command symbols used to perform control functions... ..154, 156. The output of the multiplexer 154 connects to the X port. (The interface unit 24b connects the output of the multiplexer 154 to the Y port.) The multiplexer 156... ..link 34x)) to the checker logic 160 of the packet transmitter 94y (of the interface unit 24b).

A selection (S) input of the multiplexers receives a 1-bit output from an... ..is accessible to the MP 18 via an OLAP (not shown) formed in the interface unit 24, and is written with information that "personalizes," among other things, the interface units 24. Here, the X/Y stage of the configuration register 162 configures the packet transmitter 94x of the X interface unit 24a to communicate the X encoder 150x output to the X port; the output of...traffic is present, the operation of the two packet interfaces 94 (and, thereby, the interface units 24 with which they are associated) are continually monitored. Should one of the checkers detect... ..will be asserted, resulting in an internal interrupt being posted for appropriate action by the processors 20.

Message packet traffic operates in the same manner. Assume, for the moment, that the... ..that information, a byte at a time, to the X encoder 150x of both interface units 96, which will translate each byte to encoded 9-bit form. The output of the... ..is checked with that from the packet transmitter 94x. Again, the operation of the interface units 24a, 24b, and the packet transmitters they contain, are inspected for error.

In the same... ..monitored.

Returning for the moment to Fig. 5, if the outgoing message packet is a processor initiated transaction (e.g., a read request), the processors 20 will expect a message packet to be returned in response. Thus, when the BTE... ..will issue a timeout signal to the interrupt logic (Fig. 14A) to thereby notify the processors 20 of the absence of a response to a particular transaction (e.g., a read... ..the access, to name just a few. Also, the area of memory of the memory unit 28 desired to be accessed are identified in the message packets by virtual or I... ..virtual addresses be translated to physical addresses of the memory 28. Finally, interrupts generated by units or elements external to the CPU 12A, are transmitted via message packets to interrupt the processors 20, which are also written to memory 28 when received. All this is handled by the interrupt logic and AVT logic 86, 90.

The AVT logic unit 90 utilizes a table (maintained by the processor 20 in memory 28) containing AVT entries for each possible external source permitted access to the memory 28. Each AVT entry identifies a specific source element or unit and the particular page (a page being nominally 4K (4096) bytes), or portion of a...expected" memory accesses. Expected memory accesses are those initiated by the CPU 12 (i.e., processors 20) such as a read request for information from an I/O device. These latter memory accesses are handled by a transaction sequence number (TSN) assigned to each processor initiated request. At about the time the read request is generated, the processors 20 will allocate an area of memory for the data expected to be received in... ..and 26b are, in turn, respectively coupled to the memory interfaces 70 of each interface unit 24a, 24b. The 64-bit doublewords are written to the memory 28 with the upper... ..check bits respectively from the memory interfaces 70 (70a, 70b) of each of the interface units 24a, 24b (Fig. 5).

Referring to Fig. 10, each memory interface 70 receives, from either the bus 82 from the processor interface 60 or the bus 83 from AVT logic 90 (see Fig. 5), of the associated interface unit 24, 64 bits of data to be written to memory. The busses 76 and 83... ..other for cross-checking between them. Thus, for example, the memory interface 70a (of interface unit 24a) will drive the MC 26a with the "upper" 32 bits of the 64 bits... ..are check bits, leaving 40 bits unused.

Access Validation:

As previously indicated, components of the processing system 10 external to the CPU 12A (e.g., devices of the I/O packet... ..not without qualification. Access validation, as implemented by the AVT logic 90 of the interface units 24, operates to prevent the content of the memory 28 from being corrupted by erroneously 28 are validated by the AVT logic 90 of each interface unit 24 (Fig. 5), using all of six checks: (1) that the CRC of the message... ..also are permitted the particular message packet source.

The access validation mechanism of the interface unit 24a, AVT logic 88, is shown in greater detail in Fig. 11. Incoming message packets processor 20.

The mask operation permits the size of the table of AVT entries to be varied. The content- of the AVT mask register 175 is accessible to the processor 20, permitting the processors 20 to optionally select the size of the AVT entry table. A maximum AVT table... 172 allows the AVT size to be matched to the needs of the system. A processing system 10 that includes a larger number of external elements (e.g., the number of... amount of the memory space of memory 28 to the AVT entries. Conversely, a smaller processing system 10, with a smaller number of external elements will not have such a large... set to a logic "ZERO" indicate an nonexistent TNet address, outside the limits of the processing system 10. A received packet with a TNet address outside the allowable TNet range will... in Fig. 11 as being held in the AVT entry register 180 during the validation process. AVT entries have two basic formats: normal and interrupt. The format of a normal AVT... of the AVT input register 170) will result in an error being posted to the processor via an interrupt.

A 12-bit "Permissions" field is included in t AVT entry to...path=0). Denials are logged as interrupts with the interrupt logic, and reported to the processor 20 - if the E field is set to a state ("ONE") that enables error-reporting... e.g., to a "ONE"), the other fields (Upper Bound, etc.) gain new definitions for processing interrupt writes and managing interrupt queues. This is discussed in more detail below in connection... memory 28 will be handled. Set to one state, the requested write operation will be processed normally; set to a second state, write requests specifying addresses with a fractional cache line... be written to a specific queue (interrupt queue) in memory 28, with signalling provided the processors 20 to indicate that an interrupt has been received and "posted," and ready for servicing by the processors 20. Since the interrupt queues are at specific memory locations, the processor can obtain the interrupt data when needed.

An AVT interrupt entry for an interrupt may... by the interrupt logic 86, and extracted from the head of the queue by the processor 20 when servicing the interrupt.

The AVT interrupt entry also includes a 20-bit segment ("Source ID") containing source ID information, identifying the external unit seeking attention by the interrupt process. If the source ID information of the AVT interrupt entry does not match that contained... class" of the interrupt that is used to determine the interrupt level set in the processor 20 (described more fully below); (2) a queue number that is used to select, as... capability to deliver interrupts to a CPU 12 for servicing. For example, an I/O unit may be unable to complete a read or write transaction issued by a CPU because... identify the recipient. These and other errors, exceptions, and irregularities, noted by the I/O units, or the I/O Interface elements, can become the a condition that requires the intervention... the AVT entry register 180 for use by the interrupt logic 86 of the interface

unit 24 (Fig. 5), illustrated in greater detail in Fig. 14A.

It is interrupt logic 86... four circular queues specified by the base address information contained in the AVT entry. The processor(s) 20 will then be notified, and it will be up to them as to... selected tail queue register 256 by combiner circuit 270, the output of which is the processed by the "mod z" circuit 273 to turn new offset into the queue at which... signal. The Queue Full warning signal becomes an "intrinsic" interrupt that is conveyed to the processor units 20 as a warning that if the matter is not promptly handled, later-received interrupt... will be discarded.

Incoming message packet interrupts will cause interrupts to be posted to the processor 20 by first setting one of a number of bit positions of an interrupt register 280. Multi-entry queued interrupts are set in interrupt registers 280a for posting to the processor 20; single-entry queue interrupts use interrupt register 280b. Which bit is set depends upon... multi-entry queued interrupts, soon after a multi-entry queued interrupt is determined, the interface unit will assert a corresponding interrupt signal (I1) that is applied to decode circuit 283. Decode... of register 280a to set, thereby providing advance information concerning the received interrupt to the processor(s) 20, i.e., (1) the type of interrupt posted, and (2) the class of... to one another by a compare circuit 279. The update register is writable by the processor 20 to select a register pair for comparison. If the content of the two selected... cleared.

Digressing for the moment, there are two basic types of interrupts that concern the processors 20: those interrupts that are communicated to the CPU 12 by message packets, and those...the seven interrupt postings to a latch 288, from which they are coupled to the processor 20 (20a,20b) which has an interrupt register for receiving holding the postings.

In addition... ..change in interrupts (either an interrupt has been serviced, and its posting deleted by the processor 20, or a new interrupt has been posted), a "CHANGE" signal will be issued to the processor interface 60 to inform it that an interrupt posting change has occurred, and that it should communicate the change to the processor 20.

Preferably, the AVT entry register 180 is configured to operate like a single line... ..such as set-associative, fully-associate, or direct-mapped, to name a few.

Coherency:

Data processing systems that use cache memory have long recognized the problem of coherency: making sure that... ..the incoming packet is permitted access are applied to a boundary crossing (Bdry Xing) check unit 219. Boundary check unit 219 also receives an indication of the size of the cache block the CPU 12... ..Len field of the header information from the AVT input register 170. The Bdry Xing unit determines if the data of the incoming packet is not aligned on a cache boundary...time an interrupt will be written to the queued interrupt register 280, to alert the processors 20 that a portion of the incoming data is located in the special queue.

In... ..not, the packet (both header and data) is written to a special queue, and the processors so notified by the intrinsic interrupt process described above. The processors may then move the data from the special queue to cache 22, and later write... ..the cache 22 and the memory 28 is preserved.

Block Transfer Engine (BTE):

Since the processor 20 is inhibited from directly communicating (i.e., sending) information to elements external to the... ..indirect method of information transmission.

The BTE 88 is the mechanism used to implement all processor initiated I/O traffic to transfer blocks of information. The BTE 88 allows creation of... ..BTE registers 300, 302 whose content is coupled to the MUX 306 (of the interface unit 24a; Fig. 5) and used to access the system memory 28 via the memory controllers... ..BTE data structure 304 in the memory 28 of the CPU 12A (Fig. 2). The processors 20 will write a data structure 304 to the memory 28 each time information is... ..begin on a quadword boundary, and the BTE registers 300, 302 are writable by the processors 20 only. When a processor does write one of the BTE registers 300, 302, it does so with a word... ..the request bit (rc0, rcl) to a clear state, which operates to initiate the BTE process, which is controlled by the BTE state machine 307.

The BTE registers 300, 302 also... ..cause (ec) bit differentiates time-outs and NAKs.

When information is being transferred by the processors 20 to an external unit, the data buffer portion 304b of the data structure 304 holds the information to be transferred. When information from an external unit is received by the processors 20, the data buffer portion 304b is the location targeted to hold the read response information.

The beginning of the data structure 304, portion 304a written by the processor 20, includes an information field (Dest), identifying the external element which will receive the packet... ..the transmitted data is to be written. This information is used by the packet transmitter unit 120 (Fig. 5) to assemble the packet in the form shown in Figs. 3-4...list (el) bit, when set, indicates the end of the chain, and halts the BTE processing.

The interrupt completion (ic) bit, when set, will cause the interface unit 24a to assert an interrupt (BTEComp) which sets a bit in the interrupt register 280... ..the chain pointer).

The interrupt time-out (it) bit, when set, will cause the interface unit 24a to assert an interrupt signal for the processor 20 if the acknowledgement of the access times-out (i.e., if the request timer... ..time), or elicits a NAK response (indicating that the target of the request could not process the request).

Finally, if the check sum (cs) bit is set, the data to be... ..containing the data from which the check sum was formed.

To sum up, when the processors 20 of the CPU 12A desire to send data to an external unit, they will write a data structure 304 to the memory 28, comprising identifier information in portion 304a of the data structure, and the data in the buffer portion 304b. The processors 20 will then determine the priority of the data and will write the BTE register... ..information, and sent.

If the data structure 304 indicates a read request (i.e., the processors 20 are seeking data from an external unit - either an I/O device or a CPU 12), the Len and Local Buffer Ptr... ..receiver 100 (Fig. 5) until the local memory write operation is executed.

Responses to a processor-generated read request to an external unit are not processed by the AVT table logic 146. Rather, when the processors 20 set up the BTE data structure, a transaction sequence number (TSN) is assigned the... ..the BTE 88, which will be an HAC type packet (Fig. 4) discussed above. The processors 20 will also include an memory address in the BTE data structure at which the... ..302, assume that the foregoing transfer of data from the CPU 12A to an external unit is of a large block of information. Accordingly, a number of data structures would be set up in memory 28 by the processors 20, each (except the last) including a chain pointer to additional data structures, the sum to be made by the processors 20. In such a case, the associated data structure 304 for such higher priority request... ..with another BTE operation descriptor.

Memory Controller:

Returning, for the moment, to Fig. 2, interface units 24a, 24b access the memory 28 via a pair of memory controllers (MC) 26a, 26b. The Mcs provide a fail-fast interface between the interface units 24 and the memory 28. The Mcs 26 provide the control logic necessary for accessing... ..in dynamic random access memory (DRAM) logic). The Mcs receive memory requests from the interface units 24, and execute reads and writes as well as providing refresh signals to the DRAMs... ..to provide a 72 bit data path between the memory array 28 and the interface units 24a, 24b, which utilize an SBC-DBD-SbD ECC scheme, where b=4, on a... ..26a, 26b to work together and simultaneously supply a 64-bit word to the interface units 24 with minimum latency, one-half of which (D0) comes from the MC 26a, and the other half (D1) comes from the other MC 26b. The interface unit 24 generate and check the ECC check bits. The ECC scheme used will not only... ..26 bus 25, as well as in internal registers.

From the viewpoint of the interface units 24, the memory 28 is accessed with two instructions: a "read N doubleword" and a... ..doubleword read or a block read format. The signal called "data valid" tells the interface units 24 two cycles ahead of time that read data is being returned or not being returned.

As indicated above, the maintenance processor (MP 18; Fig. 1A) has two means of access to the CPUs 12. One is... ..18 will write a register contained in the OLAP 285 with instructions that permit the processors 20 to build an image of a sequence of instructions in the memory that will permit them (the processors 20) to commence operation, going to I/O for example to transfer instructions and data from an external (storage) device that will complete the boot process.

The OLAP 285 is also used by the processors 20 to communicate to the MP 18 error indications. For example, if one of the interface units 24 detect a parity error in data received from the memory controller 26, it will...and address transfers on the bus 25 between the MC 26a and the corresponding interface unit 24a. The addressing and data transfers on the DRAM data bus, as well as generation... ..the CPU 12.

Packet Routing:

The message packets communicated between the various elements of the processing system 10 (e.g., CPUs 12A, 12B, and devices coupled to the I/O packet... ..First, each TNet Link L connects to an element (e.g., router 14A) of the processing system 10 via a port that has both receive and transmit capability. Each transmit port... ..cycle (i.e., each clock period) of the T(underscore)Clk so that the clock

synchronization FIFO at the receiving end of the transmission will maintain synchronization.

Clock synchronization is dependent upon the mode in which the processing system 10 is operated. If operating in the simplex mode in which the CPUs 12A... ..connect directly to the CPUs may drift with respect to each other. Conversely, when the processing system 10 operates in a duplex mode (e.g., the CPUs operate in synchronized, lock-step operation), the clocks between routers 14 and the CPUs 12 to which they... ..not necessarily phase-locked).

The flow of data packets between the various elements of the processing system 10 is controlled by command symbols, which may appear at any time, even within... ..initiated by a CPU 12, or MP 18, and promulgated to all elements of the processing system 10 by the routers 14 to communicate an event requiring software action by all... ..command symbol is used in conjunction with near frequency operation as an aid to maintaining synchronization between the two clock signals that (1) transfer each symbol to, and load it in each receiving clock synchronization FIFO, and (...symbols from the FIFO.

SLEEP: This command symbol is sent by any element of the processing system 10 to indicate that no additional packet (after the one currently being transmitted, if... ..received.

SOFT RESET (SRST): The SRST command symbol is used as a trigger during the processes ("synchronization" and "reintegration," described below) that are used to synchronize symbol transfers between the CPUs 12 and the routers 14A, 14B, and then to place... ..SYNC command symbol is sent by a router 14 to the CPU 12 of the processing system 10 (i.e., the sub-processor systems 10A/10B) to establish frequency-lock synchronization between CPUs 12 and routers 14A, 14B prior to entering duplex mode, or when in duplex mode to request synchronization, as will be discussed more fully below. The SYNC command symbol is used in conjunction... ..or duplex to simplex), among other things, as discussed further below in the section on Synchronization and Reintegration.

THIS LINK BAD (TLB): When any system element receiving a symbol from a TNet link L (e.g., a router, a CPU, or an I/O unit) notes an error when receiving a command symbol or packet, it will send a TLB... ..identical pairs of symbols that are compared to one another when pulled from the clock synchronization FIFOs..The DVRG command symbol signals the CPU 12 that a mis-compare has been noted. When received by the CPUs, a divergence detection process is entered whereby a determination is made by the CPUs which CPU may be failing... ..command symbols described above operate to control message flow between the various elements of the processing system 10 (e.g., CPUs 12, router 14, and the like), using principally the BUSY... ..particular TNet port however, an "end node" (i.e., a CPU 12 or I/O unit 17 - Fig. 1) may not assert backpressure because one of its transmit ports is backpressured... ..Improperly addressed packets are discarded by the router 14.

When a system element of the processing system 10 receives a BUSY command symbol on a TNet link L on which it...other command symbols (READY, BUSY, etc.).

Whenever a TNet port of an element of the processing system 10 detects receipt of a READY command symbol, it will terminate transmission of FILL... ..receives.

As will be seen, all elements (e.g., router 14, CPUs 12) of the processing system 10 that connect to a TNet link L for receiving transmitted symbols will receive those symbols via a clock synchronization (CS) FIFO. For example, as discussed above, the interface units 24 of CPUs 12 include all CS FIFOs 102x, 102y (illustrated in Fig. 6). The... ..depth to allow for speed matching, and the elastic FIFOs must provide sufficient depth for processing delays that

may occur between transmission of a BUSY command symbol during receipt of a... another data byte in packet B. As packet A progresses to the next router, the process would be repeated. If the router 14 displaces more data bytes than the FIFO can...irrespective of its own findings.

SLEEP Protocol:

The SLEEP protocol is initiated by a maintenance processor via a maintenance interface (an on-line access port - OLAP), described below. The SLEEP protocol... reintegrate a slice of the system 10. Routers 14 must be idle (no packets in process) in order to change modes without causing data loss or corruption. When a SLEEP command symbol is received, the receiving element of processing system 10 inhibits initiation of transmission of any new packet on the associated transmit port... The HALT command symbol provides a mechanism for quickly informing all CPUs 12 in a processing system 10 that is necessary to terminate I/O activity (i.e., message transmissions between... CPUs that receive HALT command symbols on either of their receive ports (of the interface units 24) will post an interrupt to the interrupt register 280 if the system halt interrupt... interrupt; Fig. 14A).

The CPUs 12 may be provided with the ability to disable HALT processing. Thus, for example, the configuration registers 75 of the interface units 24 can include a "halt enable register" that, when set to a predetermined state (eg., ZERO) disables HALT processing, but reporting detection of a HALT symbol as an error.

Router Architecture:

Referring now to... simplified block diagram of the router 14A is illustrated. The other routers 14 of the processing system 10 (e.g., routers 14B, 14', etc.) are of substantially identical construction and, therefore...for these ports 4, 5 are structured to operate in a frequency locked environment when a processing system 10 is set for duplex mode operation. In addition, when in duplex mode, a... 5025)) will receive the command/data symbols from the CPUs, pass them through the clock synchronization FIFOs 518 (discussed further below), and compare each symbol exiting the clock synchronization FIFOs with a gated compare circuit 517. When duplex operation is entered, a configuration register... 517 to activate the symbol by symbol comparison of the symbols emanating from the two synchronization FIFOs 518 of the router input logic 502 for the ports 4 and 5. Of... to that received, at substantially the same time, by the other port input.

To maintain synchronization in the duplex mode, the two port outputs of the router 14A that transmit to... mode, are duplicated by the routers 14, and returned to both CPUs.) The output logic units 5044), 5045)) that are coupled directly to the CPUs 12 will both receive symbols from... message packet identifies only one of the duplexed CPUs 12, e.g., CPU 12A) in synchronized fashion, presenting those symbols in substantially simultaneous fashion to the two CPUs 12. Of course, the CPUs 12 (more accurately, the associated interface units 24) receive the transmitted symbols with synchronizing FIFOs of substantially the same structure as that illustrated in Fig. 7A so that, even... from the FIFO structures by both CPUs 12 on the same instruction cycle, maintaining the synchronized, lock-step operation of the CPUs 12 required by the duplex operating mode.

As will... conjunction with configuration data written to registers contained in control logic 509 by the maintenance processor 18 (via the on-line access port 285' and serial bus 19A; see Fig. 1A... links L. The input logic 505 of each port input 502 also assists in maintaining synchronization - at least for those ports sending symbols in the near-frequency environment - by removing received... slower-receiving element receiving symbols from a faster-sending element could overload the input clock synchronization FIFO of the slower-receiving element. That is, if a slower clock is used to pull symbols from the clock synchronization FIFO put there by a faster clock, ultimately the clock synchronization FIFO will overflow.

The preferred technique employed here is to periodically insert SKIP symbols in... stream to avoid, or at least minimize, the possibility of an overflow of the clock synchronization FIFO (i.e., clock synchronization FIFO 518; Fig. 20A) of a router 14 (or CPU 12) due to a T... being slightly higher in frequency than the local clock used to pull symbols from the synchronization FIFO. Using SKIP symbols to by-pass a push (onto the FIFO) operation has

the... ..stall each time a SKIP command symbol is received so that, insofar as the clock synchronization FIFO is concerned, the transmitting clock that accompanied the SKIP symbol was missing.

Thus, logic... ..the port inputs 502 will recognize, and key off receipt of, SKIP command symbols for synchronization in the near frequency clocking environment so that nothing is pushed onto the FIFO, but... ..14, or between routers 14, or between a router 14 and an I/O interface unit 16A - Fig. 1) at a 50 Mhz rate, this allows for a worst case frequency... ..symbol by supplying FILL or IDLE symbols (which are received and pushed onto the clock synchronization FIFOs, but are not passed to the elastic FIFOs). In short, each elastic FIFO 506...received symbols are then communicated from the input register 516 and applied to a clock synchronization FIFO 518, also by the T(underscore)Clk. The clock synchronization FIFO 518 is logically the same as that illustrated in Figs. 8A and 8B, used in the interface units 24 of the CPUs 12. Here, as Fig. 20A shows, the clock synchronization FIFO 518 comprises a plurality of registers 520 that receive, in parallel, the output of... ..516. Associated with each of the registers 520 is a two-stage validity (V) bit synchronizer 522, shown in greater detail in Fig. 20B, and discussed below. The content of each... ..registers 520, together with the one-bit content of each associated two-stage validity bit synchronizer 522, are applied to a multiplexer 524, and the selected register/synchronizer pulled from the FIFO, and coupled to the elastic FIFO 506 by a pair of... ..is determined the state of the Push Select signal provided by a push pointer logic

unit 530; and, selection of which register 520 will supply its content, via the MUX 524... ..and loading of the register 520 selected by the push pointer logic 530. Similarly, the synchronization FIFO control logic 534 receives the clock signal local to the router (Rcv Clk) to... ..pointer logic 532.

Digressing for a moment, and referring to Fig. 20B, the validity bit synchronizer 522 is shown in greater detail as including a D-type flip-flop 541 with... ..530 (Fig. 20A) selects the register 520 of the FIFO with which the validity bit synchronizer is associated for receipt of the next symbol - if not a SKIP symbol.

The delay... ..Truth Table, below). The D-type flip-flop 543 acts as an additional stage of synchronization, ensuring a stable level at the V output relative to the local Rec Clk. The... ..flip-flop 542, allowing the Pull signal (a periodic pulse from the sync FIFO Control unit 534) to clear the validity bit on this validity synchronizer 522 when the associated register 520 has been read.

In summary, the validity synchronizer 522 operates to assert a "valid" (V) signal when a symbol is loaded in a...blocked from being routed out a particular port because another message is already in the process of being routed out that port. However, that other message in turn is also blocked... ..an incoming message packet bound for the CPUs will be replicated by the crossbar logic unit by routing the message packet to both port output 5044)) and 5045)) at the same...P) identifies which of path (X or Y) should be used for accessing two sub-processing the device.

The routers 14 provide a capability of constructing a large, versatile routing network for, for example, massively parallel processing architectures. Routers are configured according to their location (i.e., level) in the network by...expansion registers 509j)) and 509k)) are such that bits "def" are used in the algorithmic process, then bits "abc" of the Region ID are compared to the content of the Device... ..content of the route to default register 509f))) to the final stage of the selection process: check logic 602. Check logic 602 operates to check the status of the port output... ..a lower level router, and may be located in one or another of the sub-processing systems 10A, 10B. Whether a router is an upper level or lower level router depends...of CPUs 12 and I/O devices 16 to one another, forming a massively parallel processing (MPP) system. Other such MPP systems may exist, and it is those routers configured as... ..captured. As soon as the message packet's Destination ID is so captured, the selection process begins, proceeding to the development of a target port address that will be used to...an error that will be posted to the MP18 via the router's (or interface unit's) OLAP for action.

Digressing, it should be appreciated that these protocol rules observed by the routers 14 are also observed by the CPUs 12 (i.e., interface units 24) and I/O packet interfaces 17.

Finally, when the router 14A is in the... directly with the CPUs 12A, 12B, and duplex mode is used, a duplex operation logic unit 638 is utilized to coordinate the port output connected to one of the CPUs 12A... was able to write instructions to the OLAP 285 that would be executed by the processors 20 to build a small memory image and routine to permit the CPU 12 to... the clock generation circuit design. There will be one clock generator circuit in each sub-processor system 10A/10B (Fig. 1) to maintain synchronism. Designated generally with the reference numeral 650... used by the various elements (e.g. CPU. 12, routers 14, etc.) of the sub-processor system containing the clock circuit 650 (e.g., 10A).

The clock generator 654 is shown... The 50 Mhz clock signals produced by the counter 663 are distributed throughout the sub-processor system where needed.

Turning now to Fig. 25, there is illustrated the interconnection and use... the clock circuits 650 used to develop synchronous clock signals for a pair of sub-processor systems 10A, 10B (Fig. 1) for frequency locked operation. As illustrated in Fig. 25, the two CPUs 12A and 12B of the sub-processor systems 10A, 10B each have a clock circuit 650, shown in Fig. 25 as clock... 654B of both CPUs 12. A driver and signal line 667 interconnects the two sub-processor systems to deliver the M(underscore)CLK signal developed by the oscillator circuit 652A to the clock generator 654B of the sub-processor system 10B. For fault isolation, and to maintain signal quality, the M(underscore)CLK signal is delivered to the clock generator 654A of the sub-processor system 10A through a separate driver and a loopback connection 668. The reason for the... the cable (not shown) will establish the connection shown in Fig. 25 between the sub-processor systems 10A, 10B; connected another way, ...Fig. 25, the M(underscore)CLK signal produced by the oscillator circuit 652A of sub-processing system 10A is used by both sub-processing systems 10A, 10B as their respective SYNC CLK signals and the various other clock signals... produced by the clock generators 654A, 654B. Thereby, the clock signals of the paired sub-processing systems 10A, 10B are synchronized for the frequency locked operation necessary for duplex mode.

The VCXOs 662 of the clock... This allows both clock generators 654A, 654B to continue to provide to the two sub-processing systems 10A, 10B clock signals in the face of improper operation of the oscillator circuit 652A, although the sub-processor systems may no longer be frequency-locked.

The LOCK signals asserted by the phase comparators... LOCK signal signifies that the 50 Mhz signals produced by a clock generator 654 are synchronized, both in phase and in frequency, to the M(underscore)CLK signal. Thus, if either... signal that accompanies the symbol stream, and is used to push symbols onto the clock synchronizing FIFO of the receiving element (router 14, or CPU 12) is substantially identical in frequency... not phase, to that of the receiving element used to pull symbols from the clock synchronization FIFOs. For example, referring to Fig. 23, which illustrates symbols being sent from the router... clock (Local Clk). The former (Rcv Clk) is used to push symbols onto the clock synchronization FIFOs 126 of each CPU, whereas the latter is used to pull symbols from the... much higher frequency clock signal. In such situations provision must be made to ensure that synchronization is maintained between the two CPUs as to symbols pulled from the clock synchronization FIFOs 126 of each.

Here, a constant ratio clocking mechanism is used to control operation of the two clock synchronization FIFOs 126, providing the clock signal that pulls symbols from the two FIFOs at the... control mechanism is shown, designated with the reference numeral 70. As Fig. 26A illustrates, clock synchronization FIFO control mechanism 700 includes an pre-settable, multi-stage serial shift register 702, the... ratio of the clock signal at which symbols are communicated and pushed onto the clock synchronization FIFOs 126 to the frequency of the clock signal used locally. Here, a 15 stages... that will be used as the Local Clk signal to pull symbols from the clock synchronization FIFOs 126, and to operate (update) the pull pointer counter 130. The selected output is... of the CPU 12 to the clock signal used to push symbols onto the clock synchronization FIFO 126, Rev Clk, the serial shift

register is preset so that M stages of... ..duplexed CPUs 12 with a 50 Mhz clock. Thus, symbols are pushed onto the clock synchronization FIFOs 126 of the CPUs at a 50 Mhz rate. Assume further that the clock... ..of the MUX 704, which produces the clock signal that pulls symbols from the clock synchronization FIFOs 126, Rcv Clk, will contain, for each 100 ns period, five clock pulses. Thus... ..five symbols will be pushed onto, and five symbols will be pulled from, the clock synchronization FIFOs 126.

This example is symbolically shown in Fig. 26B, while the timing diagram shown... ..labelled "IN" in Fig. 27) of the Rcv Clk will push symbols onto the clock synchronization FIFOs 126. During that same 100 ns period, the serial shift register 702 circulates a... ..clocks which would require additional storage (i.e., an increase in the size of the synchronization FIFO) and impose more latency.

The constant ratio clock circuit presented here (Figs. 26) is... ..frequency to a clock regime of a different, higher frequency. The use of a clock synchronization FIFO is necessary here for compensating effects of signal delays when operating in synchronized, duplexed mode to receive pairs of identical command/data symbols from two different sources. However... ..so long as there are at least two registers in the place of the clock synchronization FIFO. Transferring data from a higher-frequency clock regime to a lower frequency clock regime... ..a wide range of possible clock ratios.

I/O Packet Interface:

Each of the sub-processor systems 10A, 10B, etc. will have some input/output capability, implemented with various peripheral units, although it is conceivable that the I/O of other sub-processor systems would be available so that a sub-processing system may not necessarily have local I/O. In any event, if local I/O...device (e.g., a signal line) would be received by the I/O packet interface unit 16 and used to form an interrupt packet that is sent to the CPU 12... ..OLAP bus, configuration information.

On-Line Access Port:

The MP 18 connects to the interface unit 24, memory controller (MC) 26, routers 14, and I/O packet interfaces with interface signals... ..OLAP 258 is essentially the same, regardless of what element (e.g. router 14, interface unit 24, etc.) it is used with. Fig. 28 diagrammatically illustrates the general structure of the... ..circuit chip used to implement certain of the elements discussed herein. For example, each interface

unit 24, memory controller 26, and router 14 is implemented by an application specific integrated circuit... ..of the OLAP 158 shown in Fig. 28 describes the OLAP associated with the interface unit 24, the MC 26, and the router 14 of the system.

As Fig. 28 shows...asymmetric variables, a "soft-vote" (SV) logic element 900 (Fig. 30A) is provided each interface unit 24 of each CPU 12. As Fig. 30 illustrates, the SV logic elements 900 of each interface unit 24 are connected to one another by a 2-bit SV bus 902, comprising bus lines 902a and 902b. Bus lines 902a carry one-bit values from the interface units 24 of CPU 12A to those of CPU 12B. Conversely, bus line 902b carries one... ..the CPU 12A.

Illustrated in Fig. 30B, is the SV logic element 900a of interface unit 24a of CPU 12A. Each SV logic element 900 is substantially identical in construction and... ..900a should be understood as applying equally to the other logic elements 900a (of interface unit 24b, CPU 12A), and 900b (of the interface units 24a, 24b of CPU 12B) unless noted otherwise. As Fig. 30B illustrates, the SV logic... ..the logic elements 900a (as well as its own). In this manner the two interface units 24a, 24b of the CPU 12A can communicate asymmetrical variables to each other.

In a... ..to the remote register 907 of logic element 902a (and that of the other interface unit 24b).

The logic elements 902 form a part of the configuration registers 74 (Fig. 5). Thus, they may be written by the processor unit(s) 20 by communicating the necessary data/address information over at least a portion of... local and remote registers 906 and 907.

The MUX 914 operates to provide each interface unit 24 of CPU 12A with selective use of the bus line 902a for the SV logic elements 900a, or for communicating a BUS ERROR signal if encountered during the reintegration process (described below) used to bring a pair of CPUs 12 into lock-step, duplex operation... same time, write the enable registers 912 of the logic element 900 of both interface units 24 of each CPU. One of the two logic elements 900 of each CPU will... it is the output enable registers 912 associated with the logic elements 900 of interface units 24a of both CPUs 12A, 12B that are written to enable the associated drivers 916. Thus, the output registers 904 of the interface units 24a of each CPU will be communicated to the bus lines 902; that is, the... to the bus line 902a, while the output register associated with logic element 900b, interface unit 24a of CPU 12B is communicated to bus line 902b. The CPUs 12 will both... again written by each CPU, followed again by reading the remote input registers 907. This process is repeated, one bit at a time, until the entire variable is communicated from the... each CPU 12 to the remote input register of the other. Note that both interface units 24 of CPU 12B will receive the bit of asymmetric information.

One example of use... elements 900 are also used to communicate bus errors that may occur during the reintegration process to be described. When reintegration is being conducted, a REINT signal will be asserted. As... ERROR signal is selected by the MUX 914 and communicated to the bus line 902a.

Synchronization:

Proper operation of the sub-processing systems 10A, 10B (Figs. 1A, 2) whether operating independently (simplex mode), or paired and operating in synchronized lock-step (duplex mode), requires assurance that data communicated between the CPUs 12A, 12B and... routers 14A, 14B will be received properly, and that any initial content of the clock synchronization FIFOs 102 (of CPUs 12A, 12B; Fig. 5) and 519 (of routers 14A, 14B; Fig... erroneously interpreted as data or commands. The push and pull pointers of the various clock synchronization FIFOs 102 (in the CPUs 12) and 518 (in the routers 14) need to be... apart, and presetting the associated FIFO queues to some known state. This done, all clock synchronization FIFOs are initialized for near frequency operation. Thus, when the system 10 is initially brought... in order to properly implement the lock-step operation of duplex mode operation, the clock synchronization FIFOs must be synchronized to operate with the particular source from which they receive data in order accommodate any... 14A, 14B to the CPUs 12A, 12B must be accounted for. It is the clock synchronization FIFOs 102 of the paired CPUs 12 that operate to receive message packet symbols, adjust... and present symbols to the two CPUs in a simultaneous manner to maintain lock-step synchronization necessary for duplex mode operation.

In similar fashion, each symbol received by the routers 14A... the CPUs (which is discussed further hereinafter). Again, it is the function of the clock synchronization FIFOs 518 of the routers 14A, 14B that receive message packets from the CPUs 12 retrieved from the clock synchronization FIFOs simultaneously.

Before discussing how the clock synchronization FIFOs of the CPUs and routers are reset, initialized, and synchronized, an understanding of their operation to maintain synchronous lock-step duplex mode operation is believed helpful. Thus, referring for the moment to Fig. 23, the clock synchronization FIFOs 102 of the CPUs 12A, 12B that receive data, for example, from the router... underscore)Clk, from the router 14A to the CPU 12B.

Consider operation of the clock synchronization FIFOs 102x)), 102y)), to receive identical symbol streams during duplex operation. Table 6, below, illustrates... held by the push and pull pointer counters 128, 130 for the CPU 12A (interface unit 24a), and the content of each of the four storage locations (byte 0. ..., byte 3... of Table 6 show the same thing for the FIFO 102y)) of CPU 12B interface unit 24a for each symbol of the duplicated symbol stream.

Assuming the delay 640 is no... ..0" locations of the queues 126. This is because (1) the FIFOs 102 have been synchronized to operate in synchronism (a process described below), and (2) the push pointer counters 128 are clocked by the clock signal... ..of the symbol stream transmitted by the router 14A will be pulled from the clock synchronization FIFOs 102 of the CPUs 12A, 12B simultaneously, maintaining the required synchronization of received data when operating in duplex mode. In effect, the depths of the queues...order to achieve the operation just described with reference to Table 6, the reset and synchronization process shown in Fig 31A is used. The process not only initializes the clock synchronization FIFOs 102 of the CPUs 12A, 12B for duplex mode operation, but also operates to adjust the clock synchronization FIFOs 518 (Fig. 19A) of the CPU ports of each of the routers 14A, 14B for duplex operation. The reset and synchronization process uses the SYNC command symbol to initiate a time period, delineated by the SYNC CLK signal 970 (Fig. 31B), to reset and initialize the respective clock synchronization FIFOs of the CPUs 12A and 12B and routers 14A, 14B. (The SYNC CLK signal... ..It is of a lower frequency than that used to receive symbols by the clock synchronization FIFOs, T(underscore)Clk. For example, where T(underscore)Clk is approximately 50 MHz, the... ..signal is approximately 3.125 MHz.)

Turning now to Fig. 31A, the reset and initialization process begins at step 950 by switching the clock signals used by the CPUs 12A, 12B and routers 14A, 14B as the transmit (T(underscore)Clk) and the unit's local clock (Local Clk) clock signals so that they are derived from the same... ..In addition, configuration registers in the CPUs 12A, 12B (configuration registers 74 in the interface units 24) and the routers 14A, 14B (contained in control logic unit 509 of routers 14A, 14B) are set to the FreqLock state.

The following discussion involves step 952, and makes reference to the interface unit 24 (Fig.5), router 14A (Fig. 19A) and Figs. 31A and 31B. With the clock... ..otherwise be sent followed by a self-addressed message packet. Any message packet in the process of being received and retransmitted when the SLEEP command symbols are received and recognized by... ..per the destination address). The SLEEP command symbol operates to "quiece" router 14A for the synchronization process. The self-addressed message packet sent by the CPU 12A, when received back by the... ..message packet sent after the SLEEP command symbol would necessarily have to be the last processed by the router 14A.

At step 954 the CPU 12A checks to see if it... ..the router will assert a RESET signal 972 that is applied to the two clock synchronization FIFOs 518 contained in the input logic 5054)), 5055)) of the router that receive symbols directly from CPUs 12A, 12B. RESET, while asserted, will hold the two clock synchronization FIFOs 518 in a temporarily non-operating reset state with the push and pull pointer... ..As each of the CPUs 12 receive SYNC symbols are detected by the storage and processing units of the packet receivers 96 (Figs. 5 and 6) cause the RESET signal to be asserted by the packet receivers 96 (actually, storage and processing elements 110; Fig. 6) of each CPU 12. the RESET signal is applied to the... ..t4))), CPUs 12 and routers 14A, 14B de-assert the RESET signals, and the clock synchronization FIFOs of the CPUs 12A, 12, and routers 14A, 14B are released from their reset...the delay, the router 14A and CPUs 12 resume pulling data from their respective clock synchronization FIFOs and resume normal operation. The clock synchronization FIFOs of the router 14A begin pulling symbols from the queue (previously set by RESET... ..from the CPU 12A with the T(underscore)Clk will be pushed onto the clock synchronization FIFO at, for example, queue location 0 (or whatever other location pointed to by the... ..0 (or whatever other location the push pointer was set to by RESET). The clock synchronization FIFOs of the router 14A are now synchronized to accommodate whatever delay 640 may be present in one communications path, relative to the... ..and the CPUs 12A, 12B.

Similarly, at the same virtual time, operation of the clock synchronization FIFOs 102 of both CPUs 12A, 12B is resumed, synchronizing them to the router 14A. Also, the CPUs 12A, 12B quit sending the SLEEP command... ..in favor of READY symbols, and resume message packet transmission, as appropriate.

That completes the synchronization process for the router 14A. However, the process must also be performed for the router 14B. Thus, the CPU 12A returns to step... ..however, assuming that the CPUs 12A, 12B are operating in duplex mode, the method and apparatus used to detect and handle a possible error, resulting in divergence of the CPUs from... ..via a message packet destined for a peripheral device of one or the other sub-processor systems 10A, 10B. Depending upon the destination of the outgoing message packet, step 1002 will... ..router 14 will issue an ERROR signal to the router control logic 509, causing the process to move to step 1004 where the router 14 detecting divergence will transmit a DVRG...time outs to occur. A router detecting divergence (without also detecting any simple link error) buys itself time to check the CRC of the received message packet by waiting for the... ..router 14, or received, all further message packets received from the CPUs and in the process of being routed when divergence was detected, or the DVRG symbol received, will be passed...1010) contained in a one of the configuration registers 74 (Fig. 5) of the interface unit 24 of each CPU.

Returning for the moment to step 1006, the determination of which... ..local" is meant to refer to the router 14A, 14B contained in the same sub-processor system 10A, 10B as the CPU. For example, referring to Fig. 1A, router 14A is... ..bit mentioned above: the bit contained in one of the configuration registers 74 of interface unit 24 (Fig. 5) of each CPU. When set to a first state, that particular CPU...the other CPU. In response, the state machines (not shown) within the control and status unit 509 (Fig. 19A) changes the "favorite" bits described above.

A few examples may facilitate understanding... ..DVRG symbol will echo that symbol to the routers 14A, 14B, start its internal divergence process timer, and begin determination of whether to continue or terminate. Having received a TLB symbol... ..to diverge with no errors reported. This can happen only if software (running on the processors 20) uses known divergent data to alter state. For example, suppose each CPU 12 has... ..number of the CPU 12A will differ from that of the CPU 12B. If the processors use the serial number to change the ...the serial number comes after some value) or to modify the value contained in a processor register, the complete "state" of the CPUs 12 will differ. In such cases, the "asymmetrical... ..of the primary CPU simply allows one CPU, and thereby the system 10, to continue processing without software intervention.

- An error at the output of the interface unit 24 of a CPU 12 will be detected by the router 14A, 14B, depending upon... ..router 14A, 14B that connects to a CPU 12 will be detected by the interface unit 24 of the affected CPU. The CPU will send a TLB symbol to the faulty... ..possible failure and, without external intervention, and transparently to the system user, remove the failing unit (CPU 12A or 12B, or router 14A or 14B) from the system to obviate or... ..reintegration." The discussion will refer to the CPUs 12A, 12B, routers 14A, 14B, and maintenance processor 18A, 18B shown forming parts of the processing system 10 illustrated in Fig. 1A. In addition, discussion will refer to the processors 20a, 20b, the interface units 24a, 24b, and the memory controllers 26a, 26b (Fig. 2) of the CPUs 12A, 12B as single units, since that is the way they function.

Reintegration is used to place two CPUs in... ..both of the paired CPUs at virtually the same time.

The major steps in the process for changing from simplex mode operation of the one on-line CPU to duplex mode... ..greater detail by the flow diagrams of Figs. 33A - 33D, generally are:

1. Setup and synchronize the two CPUs (one on-line, the other off-line) and their connected routers to... ..the memory of the on-line CPU to the off-line CPU, maintaining a tracking process that monitors changes in the memory of the on-line CPU that have not been... ..and may need to be copied over to, the off-line CPU;
3. Setup and synchronize the CPUs to run a delayed (slave) duplex mode from the same instruction stream (lock... ..will write the predetermined registers (not shown) of the control registers 74 in the interface units 24 of CPUs 12A and 12B, to a next state (after a soft operation) in...the off-line CPU 12B.

Next, a sequence is entered (steps 1060 - 1070) that will synchronize the clock synchronization FIFOs of the CPUs 12A, 12B and routers 14A, 14B in much the same fashion... ..the same steps described above in connection with the discussion of Figs. 31A, 31B to synchronize the clock synchronization FIFOs. The on-line CPU 12A will send the

sequence of a SLEEP symbol, self-addressed message packet, and SYNC symbol which, with the SYNC CLK signal, operates to synchronize CPUs and routers. Once so synchronized, the on-line CPU 12A then, at step 1066, sends a Soft Reset (SRST) command... ..of all configuration registers and control registers (e.g., configuration registers 74 of the interface units 24) cache, and the like to memory 28 of the on-line CPU, copying the... ..time to have the system 10 off-line for reintegration. For that reason, the reintegration process is performed in a manner that allows the on-line CPU to continue executing user... ..not match that of the off-line CPU. The reason for this is that normal processing by the processor 20 of the on-line CPU can change memory content after it has been copied... ..when a memory location is written in the on-line CPU 12A during the reintegration process it is marked as "dirty;" second, all copying of memory to the off-line CPU...may, however, limit the ability to detect two-bit errors. But, since the memory copying process will last for a only relatively short period of time, this risk is believed acceptable... ..memory location in CPU 12A is made (either an incoming I/O write, or a processor write operation). The returning data (that was copied over to the off-line CPU) would... ..controller 26 (Fig. 2) of the on-line CPU to monitor memory locations in the process of being copied over to the off-line CPU 12B. The memory controller uses a... ..within the block had been written by another operation (e.g., a write by the processor 20, an I/O write, etc.), that prior write operation will flag the location in ...still must be copied over to the off-line CPU 12B.

Returning to the reintegration process, and now to Fig. 33B, the memory tracking (AtomicWrite mechanism and using ECC to mark... ..entails writing a reintegration register (not shown; one of the configuration registers 74 of interface unit 24 - Fig. 5) to cause a reintegration (REINT) signal to be asserted. The REINT signal... ..is left alone. Throughout the incremental copy operations, the normal actions of the on-line processor will mark some memory locations dirty.

Several passes of incremental copying will need to be... ..the number of successful WriteConditional operations at the end of each pass through memory, the processors 20 can determine the effect of a given pass compared to the previous pass. When the benefits drop off, the processors 20 will give up on the precopy operations. At this point the reintegration process is ready to place the two CPUs 12A, 12B into lock-step operation.

Thus, the... ..in Fig. 33C, where at step 1100, the on-line CPU 12A momentarily halts foreground processing, i.e., execution of a user application. The remaining state (e.g., configuration registers, cache, etc.) of the on-line processors 20 and its caches is then read and written to a buffer (series of memory... ..to the off-line CPU 12B, together with a "reset vector" that will direct the processor units 20 of both CPUs 12A, 12B to a reset instruction.

Next, step 1106 will quiesce... ..to ensure that the FIFOs of the routers are clear, that the FIFOs of the processor interfaces 24 are clear, and no further incoming I/O message packets are forthcoming. At... ..symbol will be received and acted upon by both CPUs 12A, 12B, to cause the processor units 20 of each CPU to jump to the location in memory 28 containing the reset... ..a subroutine that will restore the stored state of both CPUs 12A, 12B to the processor units 20, caches 22, registers, etc. The CPUs 12A, 12B will then begin executing the same... ..enabling of the ECC bit to mark dirty locations must now be disabled, since the processors are doing the same thing to the same memory. During this stage of the reintegration...encountered by CPU 12A.

Meanwhile, the bus error in the CPU 12A will cause the processor unit 20 to be forced into an error-handling routine to determine (1) the cause of... ..error was caused by an attempt to read a memory location marked dirty. Accordingly, the processor unit 20 will initiate (via the BTE 88 -- Fig. 5) the AtomicWrite mechanism to copy the... ..the SRST symbols are now received by the CPUs 12A, 12B, they will cause both processor units 20 of the CPUs to be reset to start from the same location with the... ..will periodically update, e.g., a database or audit file that is indicative of the processing of the primary CPU up to that point in time of the update. Should the... ..in error-checking redundancy to the CPU 12B, in the same manner that the individual processor units 20a, 20b of the CPU 12A provide fail-fast, fault tolerance for the CPU - when...cost system is applicable, as illustrated in Fig. 34. As shown in Fig. 34, a processing system 10' includes the CPU 12A and routers 14A, 14B structured as described above. The... ..and the CPUs are also the same.

Thus, the CPU 12B' comprises only a single processor unit 20' and associated support components, including the cache 22', interface unit (IU) 24', memory controller 26', and memory 28'. Thus, while the CPU 12A is structured in the manner shown in Fig. 2, with cache processor unit, interface unit, and memory control redundancies, approximately one-half of those components are needed to implement CPU... stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... i.e., parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of... inputs connected to receive the data output from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the... attached, without the assistance of any other CPU 12 in the system. Many prior parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O device becomes... rest of the system. Other prior systems provide access to I/O through pairs of processors so that should one of the processors fail, access to the corresponding I/O is still available through the remaining I/O... if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... multiprocessing system access to every peripheral, as done here, operates to extend the "primary"/"backup" process taught in the above-identified U.S. Patent No. 4,228,496. There, a multiple CPU system may have a primary process may running on one CPU, while a backup process resides in the background on another of the CPUs. Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be... e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system... CPU can access any I/O without the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (...an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time... indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Specification: ...controller 26 (Fig. 2) of the on-line CPU to monitor memory locations in the process of being copied over to the off-line CPU 12B. The memory controller uses a... within the block had been written by another

operation (e.g., a write by the processor 20, an I/O write, etc.), that prior write operation will flag the location in... still must be copied over to the off-line CPU 12B.

Returning to the reintegration process, and now to Fig. 33B, the memory tracking (AtomicWrite mechanism and using ECC to mark... entails writing a reintegration register (not shown; one of the configuration registers 74 of interface unit 24 - Fig. 5) to cause a reintegration (REINT) signal to be asserted. The REINT signal... is left alone. Throughout the incremental copy operations, the normal actions of the on-line processor will mark some memory locations dirty.

Several passes of incremental copying will need to be... the number of successful WriteConditional operations at the end of each pass through memory, the processors 20 can determine the effect of a given pass compared to the previous pass. When the benefits drop off, the processors 20 will give up on the precopy operations. At this point the reintegration process is ready to place the two CPUs 12A, 12B into lock-step operation.

Thus, the... in Fig. 33C, where at step 1100, the on-line CPU 12A momentarily halts foreground processing, i.e., execution of a user application. The remaining state (e.g., configuration registers, cache, etc.) of the on-line processors 20 and its caches is then read and written to a buffer (series of memory... to the off-line CPU 12B, together with a "reset vector" that will direct the processor units 20 of both CPUs 12A, 12B to a reset instruction.

Next, step 1106 will quiesce... to ensure that the FIFOs of the routers are clear, that the FIFOs of the processor interfaces 24 are clear, and no further incoming I/O message packets are forthcoming. At... symbol will be received and acted upon by both CPUs 12A, 12B, to cause the processor units 20 of each CPU to jump to the location in memory 28 containing the reset... a subroutine that will restore the stored state of both CPUs 12A, 12B to the processor units 20, caches 22, registers, etc. The CPUs 12A, 12B will then begin executing the same... enabling of the ECC bit to mark dirty locations must now be disabled, since the processors are doing the same thing to the same memory. During this stage of the reintegration... encountered by CPU 12A.

Meanwhile, the bus error in the CPU 12A will cause the processor unit 20 to be forced into an error-handling routine to determine (1) the cause of... error was caused by an attempt to read a memory location marked dirty. Accordingly, the processor unit 20 will initiate (via the BTE 88 -- Fig. 5) the AtomicWrite mechanism to copy the... the SRST symbols are now received by the CPUs 12A, 12B, they will cause both processor units 20 of the CPUs to be reset to start from the same location with the... will periodically update, e.g., a database or audit file that is indicative of the processing of the primary CPU up to that point in time of the update. Should the... in error-checking redundancy to the CPU 12B, in the same manner that the individual processor units 20a, 20b of the CPU 12A provide fail-fast, fault tolerance for the CPU - when... cost system is applicable, as illustrated in Fig. 34. As shown in Fig. 34, a processing system 10' includes the CPU 12A and routers 14A, 14B structured as described above. The... and the CPUs are also the same.

Thus, the CPU 12B' comprises only a single processor unit 20' and associated support components, including the cache 22', interface unit (IU) 24', memory controller 26', and memory 28'. Thus, while the CPU 12A is structured in the manner shown in Fig. 2, with cache processor unit, interface unit, and memory control redundancies, approximately one-half of those components are needed to implement CPU... stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... i.e., parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of... inputs connected to receive the data output from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the... attached, without the assistance of any other CPU 12 in the system. Many prior parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O

device becomes... ..rest of the system. Other prior systems provide access to I/O through pairs of processors so that should one of the processors fail, access to the corresponding I/O is still available through the remaining I/O... ..if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... ..multiprocessing system access to every peripheral , as done here, operates to extend the "primary"/"backup"

process taught in the above-identified U.S. Patent No. 4,228,496. There, a multiple CPU system may have a primary process running on one CPU, while a backup process resides in the background on another of the CPUs.

Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be... ..e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system... ..CPU can access any I/O without the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... ..four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... ..an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time... ..indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... ..When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... ..response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Claims:

1. A central processor unit, comprising:

a memory for storing instructions and data;

a pair of processors operating in lock-step synchronism with each other to execute each instruction of an instruction... ..interface elements communicating the N-bit data words from corresponding ones of the pair of processors to the memory such that the first portion of the N-bit data word from a one of the pair of processors is written to the memory by the first interface unit together with the second portion of the N-bit data from the second interface element;

the first interface unit including means for receiving comparing the second portion of the N-bit data word from the second interface unit with the second portion of the N-bit data word received from the corresponding one of the pair of processors to assert an error signal is a miscompare is detected; and

the second interface unit including means for receiving comparing the first portion of the N-bit data word from the first interface unit with the first portion of the N-bit data word received from the corresponding one of the pair of processors to assert an error signal is a miscompare is detected.

Claims: ...bits.

7. Procède pour contrôler le fonctionnement d'un premier et d'un second élément processeur de données (20a, 20b) dans un système de traitement de données comportant des premier et... ..les étapes consistant :

a communiquer l'élément de données a N bits du premier élément processeur de données a un premier élément d'interface (24a / 70a) ;

a communiquer l'élément de données a N bits du second élément processeur de données a un second élément d'interface (24b / 70b) ;

a comparer, au premier élément... ..seconde partie du mot de données a N bits recue en provenance du premier élément processeur de données (20a) pour délivrer un premier signal d'erreur lorsqu'une comparaison non concordante données a N bits recue en provenance du second élément processeur de données (20b) pour délivrer un second signal d'erreur lorsqu'une comparaison non concordante... ..partie du mot de données a N bits en provenance de la première unité a processeur, et a mémoriser, en mémoire, la seconde partie du second code d'erreur a M... ..la seconde partie du mot de données a N bits en provenance du second élément processeur de données.

6/K/19 (Item 19 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Patent Assignee:

- Compaq Computer Corporation... ;
;;

Country	Number	Kind	Date
---------	--------	------	------

Abstract ...A2

Type	Pub. Date	Kind	Text
...Compaq Computer Corporation (687797) 20555 SH 249 Houston, Texas 77070-2698 US	19		
Examination...	19		
Available Text	Language	Update	Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: ...and 08/485,053, filed concurrently herewith.

The present invention is directed generally to data processing systems, and more particularly to a multiple processing system and a reliable system area network that provides connectivity for interprocessor and input/output... ..and communications systems to general purpose high availability commercial systems. The evolution of fault tolerant computers has been well documented (see D. P. Siewiorek, R. S. Swarz, "The Theory and Practice... ..and the Jet Propulsion laboratory began to apply fault tolerance to the development of guidance computers for aerospace applications. The 1960's also saw the development of the first AT&T electronic switching systems.

The first commercial fault tolerant machines were introduced by Tandem Computers in the 1970's for use in on-line transaction processing applications (J. Bartlett, "A NonStop Kernel," in proc. Eighth Symposium on Operating System Principles, pp... ..systems were introduced in the 1980's (O. Serlin, "Fault-Tolerant Systems in Commercial Applications," Computer, pp. 19-30, August 1984). Current commercial fault tolerant systems include distributed memory multi-processors, shared-memory transaction based systems, "pair-and-spare" hardware fault tolerant systems (see R. Freiburghouse, "Making Processing Fail-safe," Mini-micro Systems, pp. 255-264, May 1982; U.S. Patent No. 4... ..system.), and triple-modular-redundant systems such as the "Integrity" computing system manufactured by Tandem Computers Incorporated of Cupertino, California, assignee of this application and the invention disclosed herein.

Most applications of commercial fault tolerant computers fall into the category of on-line transaction processing. Financial institutions require high availability for electronic funds transfer, control of automatic teller machines, and... ..telecommunications systems.

Vendors of fault tolerant machines attempt to achieve both increased system availability, continuous processing, and correctness of data even in the presence of faults. Depending upon the particular system architecture, application software ("processes") running on the system either continue to run despite failures, or the processes are automatically restarted from a recent checkpoint when a fault is encountered. Some fault tolerant systems are provided with sufficient component redundancy to be able reconfigure around failed components, but processes running in the failed modules are lost. Vendors of commercial fault tolerant systems have extended fault tolerance beyond the processors and disks. To make large improvements in reliability, all sources of failure must be addressed... ..power supplies, fans and inter-module connections.

The "NonStop," and "Integrity" architectures manufactured by Tandem Computers Incorporated, (both respectively illustrated broadly in U.S. Patent No. 4,228,496 and U... ..assigned to the assignee of this application; NonStop and Integrity are registered trademarks of Tandem Computers Incorporated) represent two current approaches to commercial fault tolerant computing. The NonStop system, as generally... ..above-identified U.S. Patent No. 4,278,496, employs an architecture that uses multiple processor systems designed to continue operation despite the failure of any single hardware component. In normal operation, each processor system uses its major components independently and concurrently, rather than as "hot backups". The NonStop system architecture may consist of up to 16 processor systems interconnected by a bus for interprocessor communication. Each processor system has its own memory which contains a copy of a message-based operating system. Each processor system controls one or more input/output (I/O) busses. Dual-porting of I/O controllers and devices provides multiple paths to each device. External storage (to the processor system), such as disk storage, may be mirrored to maintain redundant permanent data storage.

This... ..hardware, while fault recovery is the responsibility of the software.

Also, in the Nonstop multi-processor architecture, application software ("process") may run on the system under the operating system as "process-pairs," including a primary process and a backup process. The primary process runs on one of the multiple processors while the backup process runs on a different processor. The backup process is usually dormant, but periodically updates its state in response to checkpoint messages from the primary process. The content of a checkpoint message can take the form of complete state update, or currently most application code runs

under transaction processing software which provides recovery through a combination of checkpoints and transaction two-phase commit protocols.

Interprocessor message traffic in the Tandem Nonstop architecture includes each processor periodically broadcasting an "I'm Alive" message for receipt by all the processors of the system, including itself, informing the other processors that the broadcasting processor is still functioning. When a processor fails, that failure will be announced and identified by the absence of the failed processor's periodic "I'm Alive" message. In response, the operating system will direct the appropriate backup processes to begin primary execution from the last checkpoint. New backup processes may be started in another processor, or the process may be run with no backup until the hardware has been repaired. U.S. Patent... ..example of this technique.

Each I/O controller is managed by one of the two processors to which it is attached. Management of the controller is periodically switched between the processors. If the managing processor fails, ownership of the controller is automatically switched to the other processor. If the controller fails, access to the data is maintained through another controller.

In addition to providing hardware fault tolerance, the processor pairs of the above-described architecture provide some measure of software fault tolerance. When a processor fails due to a software error, the backup processor frequently is able to successfully continue processing without encountering the same error. The software environment in the backup processor typically has different queue lengths, table sizes, and process mixes. Since most of the software bugs escaping the software quality assurance tests involve infrequent data dependent boundary conditions, the backup processes often succeed.

In contrast to the above-described architecture, the Integrity system illustrates another approach... ..fault recovery is the logical choice since few modifications to the software are required. The processors and local memories are configured using triple-modular-redundancy (TMR). All processors run the same code stream, but clocking of each module is independent to provide tolerance... ..three streams is asynchronous, and may drift several clock periods apart. The streams are re-synchronized periodically and during access of global memory. Voters on the TMR Controller boards detect and mask failures in a processor module. Memory is partitioned between the local memory on the triplicated processor boards and the global memory on the duplicated TMRC boards. The duplicated portions of the... ..techniques to detect failures. Each global memory is dual ported and is interfaced to the processors as well to the I/O Processors (IOPs). Standard VME peripheral controllers are interfaced to a pair of busses through a Bus... ..the BIMs to switch control of all controllers to the remaining IOP. Mirrored disk storage units may be attached to two different VME controllers. In the Integrity system all hardware failures... ..redundant hardware. After repair, components are reintegrated on-line.

US 4453215 describes a fault tolerant computer system including a plurality of CPUs, each CPU having a pair of processor sections for the purposes of fault detection. Each CPU is coupled to transfer signals onto... ..compare is detected.

The preceding examples illustrate present approaches to incorporating fault tolerance into data processing systems. Approaches involving software recovery require less redundant hardware, and offer the potential for some... ..have been developed on other systems.

Thus, the systems described above provide fault tolerant data processing either by hardware (e.g, fail-functional, employing redundancy) or by software techniques (fail-fast... ..hardware). However, none of the systems described are believed capable of providing fault tolerant data processing, using both hardware (fail-functional) and software (fail-fast) approaches, by a single data processing system.

Computing systems, such as those described above, are often used for electronic commerce: electronic data interchange (EDI) and global messaging. Today's demands upon such electronic commerce, however, is demanding

more and more throughput capacity as the number of users increases and... networks such as local area networks (LANS), and the like.

A key requirement for a server architecture is the ability to move massive quantities of data. The server should have high bandwidth that is scalable, so that added throughput capacity can be added as defined by the appended claims provides a multiple-processor system that combines both of the two above-described approaches to fault tolerant architecture, hardware redundancy and software recovery techniques, in a single system.

Broadly, the present invention includes a processing system composed of multiple sub-processing systems. Each sub-processing system has, as the main processing element, a central processing unit (CPU) that in turn comprises a pair of processors operating in lock-step, synchronized fashion to execute each instruction of an instruction stream at the same time. Each of the sub-

processing systems further include an input/output (I/O) system area network system that provides redundant communication paths between various components of the larger processing system, including a CPU and assorted peripheral devices (e.g., mass storage units, printers, and the like) of a sub-processing system, as well as between the sub-processors that may make up the larger overall processing system. Communication between any component of the processing system (e.g., a CPU and another CPU, or a CPU and any peripheral device, regardless of which sub-processing system it may belong to) is implemented by forming and transmitting packetized messages that are... responsible for choosing the proper or available communication paths from a transmitting component of the processing system to a destination component based upon information contained in the message packet. Thus, the... peripherals, but permits it to also be used for interprocessor communications.

As indicated above, the processing system of the present invention is structured to provide fault-tolerant operation through both "fail... at a variety of points in the various data paths between the (lock-step operated) processor elements of the CPU and its associated memory. In particular, the processing system of the present invention conducts error-checking at an interface, and in a manner... little impact on performance. Prior art systems typically implement error-checking by running pairs of processors, and checking (comparing) the data and instruction flow between the processors and a cache memory. This technique of error-checking tended to add delay to the... error-checking precluded use of off-the-shelf parts that may be available (i.e., processor/cache memory combinations on a single semiconductor chip or module). The present invention performs error-checking of the processors at points that operate at slower rates, such as the main memory and I/O interfaces which operate at slower speeds than the processor-cache interface. In addition, the error-checking is performed at locations that allow detection of errors that may occur in the processors, their cache memory, and the I/O and memory interfaces. This allows simpler designs for... other data integrity checks.

Error-checking of the communication flow between the components of the processing system is achieved by adding a cyclic-redundancy-check (CRC) to the message packets that... Good" (TPG) or "This Packet Bad" (TPB) - is appended to every packet. A maintenance diagnostic processor can use this information to isolate a link or router element that introduces an error... of topologies, so that alternate paths can be provided between any two elements of a processing system (e.g., between a CPU and an I/O device), for communication in the... so (e.g., by creating a "deadlock" condition, discussed further below).

The CPUs of a processing system are capable of operating in one of two basic modes: a "simplex mode" in... independently of the other, or a "duplex" mode in which pairs of CPUs operate in synchronized, lock-step fashion. Simplex mode operation provides the capability of recovering from faults ...U.S. Pat. No. 4,228,496 which teaches a multiprocessing system in which each processor has the capability of checking on the operability of its sibling processors, and of taking over the processing of a processor found or believed to have failed). When operating in duplex mode, the paired CPUs both... fault tolerant platform for less robust operating systems (e.g., the UNIX

operating system). The processing system of the present invention, with the paired, lock-step CPUs, is structured so that... ..masked (i.e., operating despite the existence of a fault), primarily through hardware.

When the processing system is operating in duplex mode, each CPU pair uses the I/O system to access any peripheral of the processing system, regardless of which (of the two, or more) sub-processor system the peripheral may be ostensibly a member of. Also, in duplex mode, message packets... ..message for the CPU pair (from either a peripheral device such as a mass storage unit or from a processing unit), will replicate the message and deliver it to both CPUs of the pair using synchronization methods that ensure that the CPUs remain synchronized. In effect, the duplex CPU pair, as viewed from the I/O system and other... ..as a single CPU. Thus, the I/O system, which includes elements from all sub-processing systems, is made to be seen by the duplex CPU pair as one homogeneous system... ..a multiprocessor system in which the CPU of any one is actually a pair of synchronized, lock-step CPUs.

Yet another important aspect of the present invention is that interrupts issuing... ..interrupts via the message packet system ensures that they will arrive at duplexed CPUs in synchronized fashion, in the same manner as I/O message packets. Interrupt message packets will contain... ..the system. In addition, using the same messaging system to communicate data between I/O units and the CPUs and to communicate interrupts to the CPUs preserves the ordering of I... ..the implementation of a technique of validating access to the memory of any CPU. The processing system permits the memory of any CPU to be accessed by any other element of...to handle input/output information transfers between a CPU and any other component of the processor system. Thereby, the individual processor units of the CPU are removed from the more mundane tasks of getting information from memory and out onto the TNet network, or accepting information from the network. The processor unit of the CPU merely sets up data structures in memory containing the data to be... ..is required, where in memory the response is to be placed when received. When the processor unit completes the task of creating the data structure, the block transfer engine is notified to... ..response is received, it is routed to the expected memory location identified, and notifies the processor unit that the response was received.

Further aspects and features of the present invention will become... ..invention, which should be taken in conjunction with the accompanying drawings.

Fig. 1A illustrates a processing system constructed in accordance with the teachings of the present invention, and Figs. 1B and 1C illustrate two alternate configurations of the processing system of Fig. 1A, employing clusters or arrangements of the processing system of Fig. 1A;

Fig. 2 illustrates, in simplified block diagram form, the central processing unit (CPU) that forms a part of each sub-processor system of Figs. 1A - 1C;

Figs. 3A - 3D and 4A - 4C illustrate the construction of... ..the area network I/O system shown in Fig. 2;

Fig. 5 illustrates the interface unit that forms a part of the CPUs of Fig. 2 to interface the processor and memory with the I/O area network system;

Fig. 6 is a block diagram, illustrating a portion of packet receiver of the interface unit of Fig. 5;

Fig. 7A diagrammatically illustrates the clock synchronization FIFO (CS FIFO) used by the packet receiver section packet receiver shown in Fig. 6;

Fig. 7B is an block diagram of a construction of the clock synchronization FIFO structure shown in Fig. 7A;

Fig. 8 illustrates the cross-connections for error-checking outbound transmissions from the two interface units of a CPU;

Fig. 9 illustrates an encoded (8B to 9B) data/command symbol;

Fig. 10 illustrates the method and structure used by the interface unit of Fig. 5 to cross-check for errors data being transferred to the memory controllers... of a CPU of Fig. 2 to other (external to the CPU) components of the processing system;

Fig. 12 is a block diagram that diagrammatically illustrates the formation of an address... 14A illustrates the logic for posting interrupt requests to queues in memory and to the processor units of the CPU of Fig. 2;

Fig. 14B illustrates the process used to form a memory address for a queue entry;

Fig. 15 is a block... data output constructs formed in the memory of the CPU of Fig. 2 by a processor unit, and containing data to be sent via the area I/O networks shown in Figs. 1A - 1C, and also illustrating the block transfer engine (BTE) unit of the interface unit of Fig. 5 that operates to access the data output constructs for transmission to the... pair of memory controllers between memory of a CPU of Fig. 2 and its interface unit for accessing from memory 72 bits of data, including two simultaneously-accessed 32-bit words... other for error-checking;

Fig. 19A is a simplified block diagram illustration of the router unit used in the area input/output networks of the processing systems shown in Figs. 1A - 1C;

Fig. 19B illustrates comparison on two port inputs of the router unit of Fig. 19A;

Fig. 20A is a block diagram the construction of one of the six input ports of the router unit shown in Fig. 19A;

Fig. 20B is a block diagram of the synchronization logic used to validate command/data symbols received at an input port of the router unit of Fig. 19A;

Fig. 21A is a block diagram illustration of the target port selection... is a block diagram illustration of one of the six output ports of the router unit shown in Fig. 19A;

Fig. 23 is an illustration of the method used to transmit identical information to a duplexed pair CPUs of Fig. 2 in synchronized fashion when the processing system is operating in lock-step (duplex) mode, using a pair the FIFOs of Fig... is a simplified block diagram illustrating the clock generation system of each of the sub-processing systems of Figs. 1A - 1C for developing the plurality of clock signals used to operate the various elements of that sub-processing system;

Fig. 25 illustrates the topology used to interconnect the clock generation systems of paired sub-processing systems for synchronizing the various clock signals of the pair of sub-processing systems to one another;

Fig. 26A and 26B illustrates a FIFO constant rate clock control logic used to control the clock synchronization FIFO of Figs. 8 or 20 in the situation when the two clocks used to... structure of the on-line access port (OLAP) used to provide access to the maintenance processor (MP) to the various elements of the system of Fig. 1A (or those of Figs... the soft-flag logic used to handle asymmetric variables between the CPUs of paired sub-processing systems operating in duplex mode;

Fig. 31A shows a flow diagram, and Fig. 31B illustrates a portion of SYNC CLK, both of which are used to reset and synchronize the clock synchronization FIFOs of the CPUs and routers of the processing system of Fig. 1A that receive information from each other;

Fig. 32 is a flow... 33A - 33D generally illustrate the procedure used to bring an one of the CPUs of processing system shown in Fig. 1A into lock-step, duplex mode operation with the other of the CPUs without measurably halting operation of the processing system; and

Fig. 34 illustrates a reduced cost architecture incorporating teachings of the invention; and... to the figures and, for the moment, principally Fig. 1A, there is illustrated a data processing system, designated with the reference 10, constructed according to the various teachings of the present invention. As Fig. 1A shows, the data processing

system 10 comprises two sub-processor systems 10A and 10B each of which are substantially the same in structure and function... should be appreciated that, unless noted otherwise, a description of any one of the sub-processor systems 10 will apply equally to any other sub-processor system 10.

Continuing with Fig. 1A therefore, each of the sub-processor systems 10A, 10B is illustrated as including a central processing unit (CPU) 12, a router 14, and a plurality of input/output (I/O) packet interfaces... one of the I/O packet interfaces 16 will also have coupled thereto a maintenance processor (MP) 18.

The MP 18 of each sub-processor system 10A, 10B connects to each of the elements of that sub-processor system via an IEEE 1149.1 test bus 17 (shown in phantom in Fig. 1A... accompanying clock signal. As Fig. 1A further illustrates, TNet Links L also interconnect the sub-processor systems 10A and 10B to one another, providing each sub-processor system 10 with access to the I/O devices of the other as well as inter-CPU communication. As will be seen, any CPU 12 of the processing system 10 can be given access to the memory of any other CPU 12, although... the memory of a CPU 12 by a wayward peripheral device 17.

Preferably, the sub-processor systems 10A/10B are paired as illustrated in Fig. 1A (and Figs 1B and 1C, discussed below), and each sub-processor system 10A/10B pair (i.e., comprising a CPU 12, at least one router 14...12A) connects, by a TNet Link L to a router (14A) of the corresponding sub-processor system (e.g., 10A). Conversely, the Y port connects the CPU (12A) to the router (14B) of the companion sub-processor system (10B). This latter connection not only provides a communication path for access by a CPU (12A) to the I/O devices of the other sub-processor system (10B), but also to the CPU (12B) of that system for inter-CPU communication.

Information is communicated between any element of the processing system 10 and any other element (e.g., CPU 12A of sub-processor system 10A) of the system and any other element of the system (e.g., an I/O device associated with an I/O packet interface 16B of sub-processor system 10B) via message "packets." Each message packet is made up of a number of... this reason, a unique method of receiving the symbols at the receiver, using a clock synchronization first-in-first-out (CS FIFO) storage structure (described more fully below), has been developed... operation means just that: the frequencies of the clock signals of the transmitter and receiver units are locked, although not necessarily in phase. Frequency locked clock signals are used to transmit symbols between the routers 14A, 14B and the CPUs 12 of paired sub-processor systems (e.g., sub-processor systems 10A, 10B, Fig. 1A). Since the clocks of the transmitting and receiving element are not phase related, a clock synchronization FIFO is again used -- albeit operating in a slightly different mode from that used for... difference, as will be seen, is due to the fact that pairs of the sub-processor systems 10 can be operated in a synchronized, lock-step mode, called duplex mode, in which each CPU 12 operates to execute the... 1A illustrates another feature of the invention: a cross-link connection between the two sub-processor systems 10A, 10B through the use of additional routers 14 (identified in Fig. 1A as... added routers RX1)), RX2)), RY1)), and RY2)) form a cross-link connection between the sub-processors 10A, 10B (or, as shown, "sides" X and Y, respectively) to couple them to I... shown in Fig. 1A, the routers RX2)) and RY2)) provide the I/O packet interface units 16x and 16y with a dual ported interface. Of course, it will now be evident... lend themselves to being used in a manner that can extend the configuration of the processing system 10 to include additional sub-processor systems such as illustrated in Figs. 1B and 1C. In Fig. 1B, for example, one... of each of the routers 14A and 14B is used to connect the corresponding sub-processor systems 10A and 10B to additional sub-processor systems 10A' and 10B' forming thereby a larger processing system comprising clusters of the basic processing system 10 of Fig. 1.

Similarly, in Fig. 1C the above concept is extended to form an eight sub-processor system cluster, comprising sub-processor systems pairs 10A/10B, 10A'/10B', 10A"/10B", and 10A'''/10B'''. In turn, each of the sub-processor systems (e.g., sub-processor system 10A) will have essentially the same basic minimum configuration of a CPU 12, a... by a I/O packet interface 16, except that, as Fig. 1C shows, the sub-processor systems 10A and 10B include additional routers 14C and 14D, respectively, in order to extend the cluster beyond sub-processor systems 10A'/10B' to the sub-processor systems 10A"/10B" and 10A'''/10B'''. As Fig. 1C further illustrates, unused ports 4 and... the routers 14 when configuring the topology of the system 10, any CPU 12 of processing system 10 of Fig. 1C can

access any other "end unit" (e.g., a CPU or I/O device) of any of the other sub-processor systems. Two paths are available from any CPU 12 to the last router 14 connecting to the I/O packet interface 16. For example, the CPU 12B of the sub-processor system 10B' can access the I/O 16''' of sub-processor system 10A''' via router 14B (of sub-processor system 10B'), router 14D, and router 14B (of sub-system 10B'') and, via link LA... 10A'''), OR via router 14A (of sub-system 10A'), router 14C, and router 14A (sub-processor system 10A''). Similarly, CPU 12A of sub-processor system 10A'' may access (via two paths) memory contained in the CPU 12B of sub-processor 10B to read or write data. (Memory accesses by one CPU 12 of another component of the processing system requires, as will be seen, the components seeking access to have authorization to do... prevents corruption of memory data of a CPU by erroneous access.).

The topology of the processing system shown in Fig. 1B is achieved by using port 1 of the routers 14A, 14B, and auxiliary TNet links LA, to connect to the routers 14A', 14B' of sub-processor systems 10A', 10B'. The topology thereby obtained establishes redundant communication paths between any CPU 12 (12A, 12B, 12A', 12B') and any I/O packet interface 16 of the processing system 10 shown in Fig. 1B. For example, the CPU 12A' of the sub-processor system 10A' may access the I/O 16A of sub-processor system 10A by a first path formed by the router 14A' (in port 4, out... shown in Fig. 1B. By interconnecting one port of each router 14 of each sub-processor pair, and using additional auxiliary TNet links LA (illustrated in Fig. 1C with the dotted line connections) between the ports 1 of the routers 14 (14A'' and 14B'') of sub-processor systems 10A'', 10B'' and 10A''', 10B''', two separate, independent data paths can be found between any CPU 12 and any I/O packet interface 16. In this fashion, any end unit (i.e., a CPU 12 or an I/O packet interface 16) will have at least two paths to any other end unit.

Providing alternate paths of access between any two end units (e.g., between a CPU 12 and any other CPU 12, or between any CPU... any two of the remaining fault domains. Here, a fault domain could be a sub-processor system (e.g., 10A). Thus, if the sub-processor system 10A were brought down because of a failure the electrical power being supplied, without... TNet link LA between the routers 14A''' and 14B''', the CPU 12B of the sub-processor system 10B would have lost access to the I/O packet interface 16''' (via router... with the loss of the router 14A (and router 14C) by loss of the sub-processor system 10A, communications between the CPU 12B is still possible via the route of router... equally to CPU 12B. As Fig. 2 shows, the CPU 12A includes a pair of processor units 20a, 20b that are configured for synchronized, lock-step operation in that both processor units 20a, 20b receive and execute identical instructions, and issue identical data and command outputs, at substantially the same moments in time. Each of the processor units 20a and 20b is connected, by a bus 21 (21a, 21b) to a corresponding cache memory 22. The particular type of processor units used could contain sufficient internal cache memory so that the cache memory 22 would not... 22 could be used to supplement any cache memory that may be internal to the processor units 20. In any event, if the cache memory 22 is used, the bus 21 is... 22 address bits, 3 bits of parity covering the address, and 7 control bits.

The processors 20a, 20b are also respectively coupled, via a separate 64-bit address/data bus 23 to X and Y interface

units 24a, 24b. If desired, the address/data communicated on each bus 23a, 23b could also be protected by parity, although this will increase the width of the bus. (Preferably, the processors 20 are constructed to include RISC R4000 type microprocessors, such as are available from the MIPS Division of Silicon Graphics, Inc. of Santa Clara, California.)

The X and Y interface units 24a, 24b operate to communicate data and command signals between the processor units 20a, 20b and a memory system of the CPU 12A, comprising a memory controller (MC... MC halves 26a and 26b) and a dynamic random access memory array 28. The interface units 24 interconnect to each other and to the Mcs 26a, 26b by a 72-bit... accompanied by 8 bits of ECC) are written to the memory 28 by the interface units 24, one interface unit 24 will drive only one word (e.g., the 32 most significant portion) of the doubleword being written while the other interface unit 24 writes the other word of the double word (e.g., the least significant 32-bit portion of

the doubleword). In addition, on each write operation the interface units 24a, 24b perform a cross-check operation on the data not written by that interface unit 24 with the data written by the other to check for errors; on read operations... ..accessed corresponds to the address of the location from which the doubleword was stored.

Interface units 24a, 24b of the CPU 12A form the circuitry to respectively service the X and Y (I/O) ports of the CPU 12A. Thus, the X interface unit 24a connects by the bi-directional TNet Link Lx to a port of the router 14A of the processor system 10A (Fig. 1A) while the Y interface unit 24b similarly connects to the router 14B of the processor system 10B by TNet Link Ly. The X interface unit 24a handles all I/O traffic between the router 14A and the CPU 12A of the sub-processor system 10A. Likewise, the Y interface unit 24b is responsible for all I/O traffic between the CPU 12A and the router 14B of companion sub-processor system 10B.

The TNet Link Lx connecting the X interface unit 24a to the router 14A (Fig. 1) comprises, as above indicated, two 10-bit buses... ..bus 32x)) carries data incoming from the router 14A. In similar fashion, the Y interface unit 24b is connected to the router 14B (of the sub-processor system 10B) by two 10-bit busses: 30y)) (for outgoing transmissions) and 32y)) (for incoming transmissions), together forming the TNet Link Ly.

The X and Y interface units 24a, 24b are synchronously operated in lock-step, performing substantially the same operations at substantially the same times. Thus, although only the X interface unit 24a actually transmits data onto the bus 30x)), the same output data is being produced by the Y interface unit 24b, and used for error-checking. The Y interface unit 24b output data is coupled to the X interface unit 24a by a cross-link 34, where it is received by the X interface unit 24a and compared against the same output data produced by the X interface unit. In this way the outgoing data made available at the X port of the CPU... ..the port of the CPU 12A is checked. The output data from the Y interface unit 24b is coupled to the Y port by a 10-bit bus 30y)), and also to the X interface unit 24a by the 9-bit cross-link 34y)) where is checked with that produced by the X interface unit.

As mentioned, the two interface units 24a, 24b operate in synchronous, lock-step with one another, each performing substantially the same... ..X and/or Y ports of the CPU 12A must be received by both interface units 24a, 24b to maintain the two interface units in this lock-step mode. Thus, data received by one interface unit 24a, 24b is passed to the other, as indicated by the dotted lines and 9... ..connections 36x)) (communicating incoming data being received at the X port by the X interface unit 24a to the Y interface unit 24b) and 36y)) (communicating data received at the Y port by the Y interface unit 24b to the X interface unit 24a).

Certain more robust operating systems are structured with a fault-tolerant capability in the... ..example, U.S. Patent No. 4,817,091 teaches a multiprocessor system in which each processor periodically messages each of the processors of the system (including itself), under software control, to thereby provide an indication of continuing operation. Each of the processors, in addition to performing its normal tasks, operates as a backup processor to another of the processors. In the event one of the backup processors fails to receive the messaged indication from a sibling processor, it will take over the operation of that sibling (now thought to be inoperative), in... ..platform for both types of software. Thus, when a robust operating system is available, the processing system 10 can be configured to operate in a "simplex" mode in which each of... ..left, in most instances, to software.

Alternatively, for less robust operating systems and software, the processing system 10 provides a hardware-based fault-tolerance by being configured to operate in a... ..g., CPUs 12A, 12B) are coupled together as shown in Fig. 1A, to operate in synchronized, lock-step fashion, executing the same instructions at the substantially the same moment in time... ..data and command symbols. In order to simplify the design of the CPU 12, the processors 20 are precluded from communicating directly with any outside entity (e.g., another CPU 12... ..O device via the I/O packet interface 16). Rather, as will be seen, the processor will construct a data structure in memory and turn over control to the interface units 24. Each interface unit 24 includes a block transfer engine (BTE; Fig. 5) configured to provide a form of... ..to the destination according to information contained in the message packet.

The design of the processing system 10 permits a memory 28 of a CPU to be read or written by... via the routers 14. Accordingly, before continuing with the description of the construction of the processing system 10, it would be of advantage to understand first the configuration of the data...information.

As indicated, the HADC message packet operates to communicate write data between the end units (e.g., CPU 12) of the processing system 10. Other message packets, however, may be differently constructed because of their function and... CRC. The HC message packet is used to acknowledge a request to write data.

Interface Unit:

The X and Y interface units 24 (i.e., 24a and 24b - Fig. 2) operate to perform three major functions within the CPU 12: to interface the processors 20 to the memory 28; to provide an I/O service that operates transparently to, but under the control of, the processors; and to validate requests for access to the memory 28 from outside sources.

Regarding first the interface function, the X and Y interface units 24a, 24b operate to respectively communicate processors 20a, 20b to the memory controllers (Mcs 26a, 26b) and memory 28 for writing and... fast checking of the data read/written. For example, write operations have the two interface units 24a, 24b cooperating to cross-check the data to be written to ensure its integrity (and at the same time, the interface units 24 will operate) to develop an error correcting code (ECC) that covers, as will be retrieved from the appropriate address.

With respect to I/O access, the processors 20 are not provided with the ability to communicate directly with the input/output systems... must write data structures to the memory 28 and then pass control to the interface units 24 which perform a direct memory access (DMA) operation to retrieve those data structures, and... indicated in the data structure itself.)

The third function of the X and Y interface units 24, access validation to the memory 28, uses an address validation and translation (AVT) table maintained by the interface units. The AVT table contains an address for each system component (e.g., an I/O... the incoming message packets are virtual addresses. These virtual addresses are translated by the interface unit to physical addresses recognizable by the memory control units 26 for accessing the memory 28.

Referring to Fig. 5, illustrated is a simplified block diagram of the X interface unit 24a of the CPU 12A. The companion Y interface unit 24b (as well as the interface units 24 of the CPU 12B, or any other CPU 12) is of substantially identical construction. Accordingly, it will be understood that a description of the interface unit 24a will apply equally to the other interface units 24 of the processing system 10.

As Fig. 5 illustrates, the X interface unit 24a includes a processor interface 60, a memory interface 70, interrupt logic 86, a block transfer engine (BTE) 88, access validation and translation logic 90, a packet transmitter 94, and a packet receiver 96.

Processor Interface:

The processor interface 60 handles the information flow (data and commands) between the processor 20a and the X interface unit 24a. A processor bus 23, including a 64 bit address and data bus (SysAD) 23a and a 9 bit command bus 23b, couples the processor 20a and the processor interface 60 to one another. While the SysAD bus 23a carries memory address and data... and qualifying commands carried at substantially the same time on the SysAD bus 23a. The processor interface 60 operates to interpret commands issued by the processor unit 20a in order to pass reads/writes to memory or control registers of the processor interface. In addition, the processor interface 60 contains temporary storage (not shown) for buffering addresses and data for access to... 26). Data and command information read from memory is similarly buffered en route to the processor unit 20a, and made available when the processor unit is ready to accept it. Further, the processor interface 60 will operate to generate the necessary interrupt signalling for the X interface unit 24a.

The processor interface 60 is connected to a memory interface 70 and to configuration registers 74 by a bi-directional 64 bit

processor address/data bus 76. The configuration registers 74 are a symbolic representation of the various control registers contained in other components of the X interface unit 24a, and will be discussed when those particular components are discussed. However, although not specifically... ..throughout other of the logic that is used to implement the X interface 24a, the processor address/data bus 76 is likewise coupled to read or write to those registers.

Configuration registers 74 are read/write accessible to the processor 20a; they allow the X interface unit to be "personalized." For example, one register identifies the node address of the CPU 12A... ..with the CPU 12A; another, readable only, contains a fixed identification number of the interface unit 24, and still other registers define areas of memory that can be used by, for... ..logic 90, etc.) employing them are discussed.

The memory interface 70 couples the X interface unit 24a to the memory controllers 26 (and to the Y interface unit 24b; see fig. 2) by a bus 25 that includes two 36 bi-directional bit... ..25a, 25b. The memory interface operates to arbitrate between requests for memory access from the processor unit 20, the BTE 88, and the AVT logic 90. In addition to memory accesses from the processor unit 20a, the memory 28 may also be accessed by components of the processing system 10 to, for example, store data requested to be read by the processor unit 20a from an I/O unit 17, or memory 28 may also be accessed for I/O data structures previously set up in memory by the processor unit. Since these accesses are all asynchronous, they must be arbitrated, and the memory interface 70... ..command information accessed from the memory 28 is coupled from the memory interface to the processor interface 60 by a memory read bus 82, as well as to an interrupt logic interface units 24a and 24b formulate and apply the (64-bit) doubleword to the bus 25, each... ..by the memory interface 70 are coupled to the memory interface by the companion interface unit 24 where they are compared with the same 32 bits for error.

Digressing for the... ..containing interrupt information are received, that information is conveyed to the interrupt logic 86 for processing and posting for action by the processor 20, along with any interrupts generated internal to the CPU 12A. Internally generated interrupts will... ..register 71 (internal to the interrupt logic 86), indicating the cause of the interrupt. The processor 20 can then read and act upon the interrupt. The interrupt logic is discussed more fully below.

The BTE 88 of the X interface unit 24a operates to perform direct memory accesses, and provides the mechanism that allows the processors 20 to access external resources. The BTE 88 can be set-up by the processors 20 to generate I/O requests, transparent to the processors 20 and notify the processors when the requests are complete. The BTE logic 88 is discussed further below.

Requests for... ..8 byte wide format necessary for storing in the memory 28.

Outgoing message packets containing processor originated transaction requests (e.g., a read request asking for a block data from an I/O unit) are monitored by the request transaction logic (RTL) 100. The RTL 100 provides a time... ..will generate an interrupt (handled and reported by the interrupt logic 86) to inform the processor 20 that the request was not honored. In addition, the RTL 100 will validate responses... ..28 (by the DMA operation of the BTE 88) at a location known to the processor 20 so that it can locate the response.

Each of the CPUs 12 are checked... ..discussed. One such check is an on-going monitor of the operation of the interface units 24a, 24b of each CPU. Since the interface units 24a, 24b operate in lock-step synchronism checking can be performed by monitoring the operating states of the paired interface units 24a, 24b by a continuous comparison of certain of their internal states. This approach is implemented by using one stage of a state machine (not shown) contained in the unit 24a of CPU 12A, and comparing each state assumed by that stage with its identical state machine stage in the interface unit 24b. All units of the interface units 24 use state machines to control their

operations. Preferably, therefore, a state machine of the memory interface 70 that controls the data transfers between the interface unit 24 and the MC 26 is used. Thus, a selected stage of the state machine used in the memory interface 70 of the interface unit 24a is selected. An identical stage of a state machine of one of the interface unit 24b is also selected. The two selected stages are communicated between the interface units 24a, 24b and received by a compare circuit contained in both interface units 24a, 24b. As the interface units operate lock-step with one another, the state machines will likewise march through the same identical states, assuming each state at substantially the same moments in time. If an interface unit encounters an error, or fails, that activity will cause the interface units to diverge, and the state machines will assume different states. The time will come when... ..that will bring to the attention of the CPUs 12A (or 12B) that the interface units 24a, 24b of that CPU are no longer in lock-step, and to act accordingly... ..X port, receiving only those message packets transmitted by the router 14A of the sub-processor system 10A (Fig. 1A). The Y port is serviced by the Y interface unit 24b to receive message packets from the router 14B of the companion sub-processor system 10B. However, both interfaces (as well as Mcs 26 and processor 20), as has been indicated, are basically mirror images of one another in that both... ..in both structure and function. For this reason, message packet information, received by one interface unit (e.g., 24a) must be passed for processing also to the companion interface unit (e.g., 24b). Further, since both interface units 24a, 24b will assemble the same message packets for transmission from the X or the Y ports, the message packet being transmitted by the interface unit (e.g., 24b) actually being communicated from the associated port (e.g., the Y port) will also be coupled to the other interface unit (e.g., 24a) for cross-checking for errors. These features are illustrated in Figs. 6... ..receiving portions of the packet receivers 96 (96x, 96y) of the X and Y interface units 24a, 24b are broadly illustrated. As shown, each packet receiver 96x, 96y has a clock... ..receive a corresponding one of the TNet Links 32. The CS FIFOs 102 operate to synchronize the incoming command/data symbols to the local clock of the packet receiver 96, buffering... ..104x, coupled to the MUX 104y of the packet receiver 96y of the Y interface unit 24b by the cross-link connection 36x)). In similar fashion, information received at the Y port is coupled to the X interface unit 24a by the cross-link connection 36y)). In this manner, the command/data symbols of... ..packets received at one of the X, Y ports by the corresponding X, Y, interface unit 24a, 24b is passed to the other so that both will process and communicate the same information on to other components of the interface units 24 and/or memory 28.

Continuing with Fig. 6, depending upon which port X, Y... ..or the other of the CS FIFOs 102x, 102y for communication to the storage and processing logic 110 of the interface unit 24. The information contained in each 9-bit symbol is an 8-bit byte of... ..the encoding of which is discussed below with respect to Fig. 9. The storage and processing logic 110 will first translate the 9-bit symbols to 8-bit data or command... ..the outputs of the CS FIFOs 102x, 102y are also coupled to a command decode unit in addition to the MUX 104. The command decode unit operates to recognize command symbols (differentiating them from data symbols in a manner that is... ..below), decoding them to generate therefrom command signals that are applied to a receiver control unit, a state machine-based element that functions to control packet receiver operations.

As indicated above... ..at the output of the MUX 104, the receiver control portion of the storage control unit enables CRC check logic 106 to calculate a CRC symbol while the data symbols are... ..below, CS FIFOs are found not only in the packet receivers 96 of the interface units 24, but also at each receiving port of the routers 14 and the I/O... ..an even more important part, and perform a unique function, when a pair of sub-processor systems are operating in duplex mode and the two CPUs 12A and 12B of the sub-processor systems 10A, 10B operate in synchronized, lock-step, executing the same instructions at the same time. When operating in this latter... ..difficult to ensure that the clocking regime of the routers 14A and 14B are exactly synchronized to those of the CPUs 12A and 12B - even when using frequency locked clocking. In... ..used to transmit symbols to a CPU 12 and the clock used by an interface unit 24 to receive those symbols.

The structure of the CS FIFO 102 is diagrammatically illustrated... ..i.e., a packet) or IDLE symbols - except during certain situations (e.g., reset, initialization, synchronization and others discussed below). As explained above, each symbol held in the transmit register 120... ..same symbol leaving the storage queue, allowing each symbol entering the storage queue 126 to settle before it is clocked out and passed to the storage and processing units 110x (and 110y)

by the MUX 104x (and 104y). Since the transmit and receive clocks... ..functioning in duplex mode) operate to transmit symbols with near frequency clocking. Even so, clock synchronization FIFOs are used at these other ports to receive symbols transmitted with near frequency clocking, and the structure of these clock synchronization FIFOs are substantially the same as that used in frequency locked environments, i.e., that... ..of the storage queue 126 are nine bits wide; in near frequency environments, the clock

synchronization FIFOs use symbol locations of the queue 126 that are 10 bits wide, the extra... ..the faster clock source. To handle this clock drift, the two pointers are effectively re-synchronized periodically.

When the CPUs 12 are paired and operating in duplex mode, all four interface units 24 operate in lock-step to, among other things, transmit the same data and receive... ..simplex mode, each independent of the other, clocking need only be near frequency.

The interface unit 24 receives a SYNC CLK signal that is used in combination with a SYNC command symbol to initialize and synchronize the Rcv register 124 to the transmitting router 14. When using either near frequency or... ..102X preferably begin from some known state. Incoming symbols are examined by the storage and processing units 110 of the packet receivers 96. The storage and processing units look for, and act upon as appropriate, command symbols. Pertinent here is that when the... ..receives a SYNC command symbol it will be decoded and detected by the storage and processing unit 110. Detection of the SYNC command symbol by the storage and processing unit 110 causes assertion of a RESET signal. The RESET signal, under synchronous control of the SYNC CLK signal, is used to reset the input buffers (including the clock synchronization buffers) to predetermined states, and synchronize them to the routers 14.

The synchronization of the CS FIFOs 102 of the interface units 24 those of one or both routers 14A, 14B is discussed more fully below in the section discussing synchronization.

Packet Transmitter:

Each interface unit 24 is assigned to transmit from and receive at only one of the X or Y ports of the CPU 12. When one of the interface units 24 transmits, the other operates to check the data being transmitted. This is an important... ..shows, in abbreviated form, the packet transmitters 94x, 94y of the X and Y interface units 24a, 24b, respectively. Both packet transmitters are identically constructed, so that discussion of one (packet... ..logic 152 that receives, from the BTE 88 or AVT 90 of the associated interface unit (here, the X interface unit 24a) the data to be transmitted - in doubleword (64-bit) format. The packet assembly logic... ..and Y ports: they are either symbols that make up a message packet in the process of being transmitted, or IDLE symbols, or other command symbols used to perform control functions... ..154, 156. The output of the multiplexer 154 connects to the X port. (The interface unit 24b connects the output of the multiplexer 154 to the Y port.) The multiplexer 156... ..link 34x)) to the checker logic 160 of the packet transmitter 94y (of the interface unit 24b).

A selection (S) input of the multiplexers receives a 1-bit output from an... ..is accessible to the MP 18 via an OLAP (not shown) formed in the interface unit 24, and is written with information that "personalizes," among other things, the interface units 24 Here, the X/Y stage of the configuration register 162 configures the packet transmitter 94x of the X interface unit 24a to communicate the X encoder 150x output to the X port; the output of... ..traffic is present, the operation of the two packet interfaces 94 (and, thereby, the interface units 24 with which they are associated) are continually monitored. Should one of the checkers detect... ..will be asserted, resulting in an internal interrupt being posted for appropriate action by the processors 20.

Message packet traffic operates in the same manner. Assume, for the moment, that the... ..that information, a byte at a time, to the X encoder 150x of both interface units 96, which will translate each byte to encoded 9-bit form. The output of the... ..is checked with that from the packet transmitter 94x. Again, the operation of the interface units 24a, 24b, and the packet transmitters they contain, are inspected for error.

In the same... ..monitored.

Returning for the moment to Fig. 5, if the outgoing message packet is a processor initiated transaction (e.g., a read request), the processors 20 will expect a message packet to be returned in response. Thus, when the BTE... ..will issue a timeout signal to the interrupt logic (Fig. 14A) to thereby notify the processors 20 of the absence of a response to a particular transaction (e.g., a read... ..the access, to name just a few. Also, the area of memory of the memory unit 28 desired to be accessed are identified in the message packets by virtual or I... ..virtual addresses be translated to physical addresses of the memory 28. Finally, interrupts generated by units or elements external to the CPU 12A, are transmitted via message packets to interrupt the processors 20, which are also written to memory 28 when received. All this is handled by the interrupt logic and AVT logic 86, 90.

The AVT logic unit 90 utilizes a table (maintained by the processor 20 in memory 28) containing AVT entries for each possible external source permitted access to the memory 28. Each AVT entry identifies a specific source element or unit and the particular page (a page being nominally 4K (4096) bytes), or portion of a... ..expected" memory accesses. Expected memory accesses are those initiated by the CPU 12 (i.e., processors 20) such as a read request for information from an I/O device. These latter memory accesses are handled by a transaction sequence number (TSN) assigned to each processor initiated request. At about the time the read request is generated, the processors 20 will allocate an area of memory for the data expected to be received in... ..and 26b are, in turn, respectively coupled to the memory interfaces 70 of each interface unit 24a, 24b. The 64-bit doublewords are written to the memory 28 with the upper... ..check bits respectively from the memory interfaces 70 (70a, 70b) of each of the interface units 24a, 24b (Fig. 5).

Referring to Fig. 10, each memory interface 70 receives, from either the bus 82 from the processor interface 60 or the bus 83 from AVT logic 90 (see Fig. 5), of the associated interface unit 24, 64 bits of data to be written to memory. The buses 82 and 83... ..other for cross-checking between them. Thus, for example, the memory interface 70a (of interface unit 24a) will drive the MC 26a with the "upper" 32 bits of the 64 bits... ..are check bits, leaving 40 bits unused.

Access Validation:

As previously indicated, components of the processing system 10 external to the CPU 12A (e.g., devices of the I/O packet... ..not without qualification. Access validation, as implemented by the AVT logic 90 of the interface units 24, operates to prevent the content of the memory 28 from being corrupted by erroneously...Accesses to the memory 28 are validated by the AVT logic 90 of each interface unit 24 (Fig. 5), using all of six checks: (1) that the CRC of the message... ..also are permitted the particular message packet source.

The access validation mechanism of the interface unit 24a, AVT logic 88, is shown in greater detail in Fig. 11. Incoming message packets...and post an interrupt to the interrupt logic 86 (Fig. 5) for action by the processor 20.

The mask operation permits the size of the table of AVT entries to be varied. The content of the AVT mask register 175 is accessible to the processor 20, permitting the processors 20 to optionally select the size of the AVT entry table. A maximum AVT table... ..172 allows the AVT size to be matched to the needs of the system. A processing system 10 that includes a larger number of external elements (e.g., the number of... ..amount of the memory space of memory 28 to the AVT entries. Conversely, a smaller processing system 10, with a smaller number of external elements will not have such a large... ..set to a logic "ZERO" indicate a nonexistent TNet address, outside the limits of the processing system 10. A received packet with a TNet address outside the allowable TNet range will... ..in Fig. 11 as being held in the AVT entry register 180 during the validation process. AVT entries have two basic formats: normal and interrupt. The format of a normal AVT...of the AVT input register 170) will result in an error being posted to the processor via an interrupt.

A 12-bit "Permissions" field is included in the AVT entry to... ..path=0). Denials are logged as interrupts with the interrupt logic, and reported to the processor 20 - if the B field is set to a state ("ONE") that enables error-

reporting... ..e.g., to a "ONE"), the other fields (Upper Bound, etc.) gain new definitions for processing interrupt writes and managing interrupt queues. This is discussed in more detail below in connection... ..memory 28 will be handled. Set to one state, the requested write operation will be processed normally; set to a second state, write requests specifying addresses with a fractional cache line...be written to a specific queue (interrupt queue) in memory 28, with signalling provided the processors 20 to indicate that an interrupt has been received and "posted," and ready for servicing by the processors 20. Since the interrupt queues are at specific memory locations, the processor can obtain the interrupt data when needed.

An AVT interrupt entry for an interrupt may... ..by the interrupt logic 86, and extracted from the head of the queue by the processor 20 when servicing the interrupt.

The AVT interrupt entry also includes a 20-bit segment ("Source ID") containing source ID information, identifying the external unit seeking attention by the interrupt process. If the source ID information of the AVT interrupt entry does not match that contained... ..class" of the interrupt that is used to determine the interrupt level set in the processor 20 (described more fully below); (2) a queue number that is used to select, as will...capability to deliver interrupts to a CPU 12 for servicing. For example, an I/O unit may be unable to complete a read or write transaction issued by a CPU because... ..identify the recipient. These and other errors, exceptions, and irregularities, noted by the I/O

units, or the I/O Interface elements, can become a condition that requires the intervention of... ..the AVT entry register 180 for use by the interrupt logic 86 of the interface unit 24 (Fig. 5), illustrated in greater detail in Fig. 14A.

It is interrupt logic 86...four circular queues specified by the base address information contained in the AVT entry. The processor(s) 20 will then be notified, and it will be up to them as to... ..selected tail queue register 256 by combiner circuit 270, the output of which is the processed by the "mod z" circuit 273 to turn new offset into the queue at which... ..signal. The Queue Full warning signal becomes an "intrinsic" interrupt that is conveyed to the processor units 20 as a warning that if the matter is not promptly handled, later-received interrupt... ..will be discarded.

Incoming message packet interrupts will cause interrupts to be posted to the processor 20 by first setting one of a number of bit positions of an interrupt register 280. Multi-entry queued interrupts are set in interrupt registers 280a for posting to the processor 20; single-entry queue interrupts use interrupt register 280b. Which bit is set depends upon... ..multi-entry queued interrupts, soon after a multi-entry queued interrupt is determined, the interface unit will assert a corresponding interrupt signal (I1) that is applied to decode circuit 283. Decode... ..of register 280a to set, thereby providing advance information concerning the received interrupt to the processor(s) 20, i.e., (1) the type of interrupt posted, and (2) the class of... ..to one another by a compare circuit 279. The update register is writable by the processor 20 to select a register pair for comparison. If the content of the two selected... ..cleared.

Digressing for the moment, there are two basic types of interrupts that concern the processors 20: those interrupts that are communicated to the CPU 12 by message packets, and those...the seven interrupt postings to a latch 288, from which they are coupled to the processor 20 (20a,20b) which has an interrupt register for receiving holding the postings.

In addition... ..change in interrupts (either an interrupt has been serviced, and its posting deleted by the processor 20, or a new interrupt has been posted), a "CHANGE" signal will be issued to the processor interface 60 to inform it that an interrupt posting change has occurred, and that it should communicate the change to the processor 20.

Preferably, the AVT entry register 180 is configured to operate like a single line... ..such as set-associative, fully-associate, or direct-mapped, to name a few.

Coherency:

Data processing systems that use cache memory have long recognized the problem of coherency: making sure that... ..the incoming packet is permitted access are applied to a boundary crossing (Bdry Xing) check unit 219. Boundary check unit 219 also receives an indication of the size of the cache block the CPU 12... ..Len field of the header information from the AVT input register 170. The Bdry Xing unit determines if the data of the incoming packet is not aligned on a cache boundary...time an interrupt will be written to the queued interrupt register 280, to alert the processors 20 that a portion of the incoming data is located in the special queue.

In... ..not, the packet (both header and data) is written to a special queue, and the processors so notified by the intrinsic interrupt process described above. The processors may then move the data from the special queue to cache 22, and later write... ..the cache 22 and the memory 28 is preserved.

Block Transfer Engine (BTE):

Since the processor 20 is inhibited from directly communicating (i.e., sending) information to elements external to the... ..indirect method of information transmission.

The BTE 88 is the mechanism used to implement all processor initiated I/O traffic to transfer blocks of information. The BTE 88 allows creation of... ..BTE registers 300, 302 whose content is coupled to the MUX 306 (of the interface unit 24a; Fig. 5) and used to access the system memory 28 via the memory controllers... ..BTE data structure 304 in the memory 28 of the CPU 12A (Fig. 2). The processors 20 will write a data structure 304 to the memory 28 each time information is... ..begin on a quadword boundary, and the BTE registers 300, 302 are writable by the processors 20 only. When a processor does write one of the BTE registers 300, 302, it does so with a word... ..the request bit (rc0, rc1) to a clear state, which operates to initiate the BTE process, which is controlled by the BTE state machine 307.

The BTE registers 300, 302 also... ..cause (ec) bit differentiates time-outs and NAKs.

When information is being transferred by the processors 20 to an external unit, the data buffer portion 304b of the data structure 304 holds the information to be transferred. When information from an external unit is received by the processors 20, the data buffer portion 304b is the location targeted to hold the read response information.

The beginning of the data structure 304, portion 304a written by the processor 20, includes an information field (Dest), identifying the external element which will receive the packet... ..the transmitted data is to be written. This information is used by the packet transmitter unit 94 (Fig. 5) to assemble the packet in the form shown in Figs. 3-4...list (el) bit, when set, indicates the end of the chain, and halts the BTE processing.

The interrupt completion (ic) bit, when set, will cause the interface unit 24a to assert an interrupt (BTEComp) which sets a bit in the interrupt register 280... ..the chain pointer).

The interrupt time-out (it) bit, when set, will cause the interface unit 24a to assert an interrupt signal for the processor 20 if the acknowledgement of the access times-out (i.e., if the request timer... ..time), or elicits a NAK response (indicating that the target of the request could not process the request).

Finally, if the check sum (cs) bit is set, the data to be... ..containing the data from which the check sum was formed.

To sum up, when the processors 20 of the CPU 12A desire to send data to an external unit, they will write a data structure 304 to the memory 28, comprising identifier information in portion 304a of the data structure, and the data in the buffer portion 304b. The processors 20 will then determine the priority of the data and will write the BTE register... ..information, and sent.

If the data structure 304 indicates a read request (i.e., the processors 20 are seeking data from an external unit - either an I/O device or a CPU 12), the Len and Local Buffer Ptr... ..receiver 100 (Fig. 5) until the local memory write operation is executed.

Responses to a processor-generated read request to an external unit are not processed by the AVT table logic 146. Rather, when the processors 20 set up the BTE data structure, a transaction sequence number (TSN) is assigned the... the BTE 88, which will be an HAC type packet (Fig. 4) discussed above. The processors 20 will also include an memory address in the BTE data structure at which the...302, assume that the foregoing transfer of data from the CPU 12A to an external unit is of a large block of information. Accordingly, a number of data structures would be set up in memory 28 by the processors 20, each (except the last) including a chain pointer to additional data structures, the sum... sent. Assume now that a higher priority request is desired to be made by the processors 20. In such a case, the associated data structure 304 for such higher priority request... with another BTE operation descriptor.

Memory Controller:

Returning, for the moment, to Fig. 2, interface units 24a, 24b access the memory 28 via a pair of memory controllers (MC) 26a, 26b. The Mcs provide a fail-fast interface between the interface units 24 and the memory 28. The Mcs 26 provide the control logic necessary for accessing... in dynamic random access memory (DRAM) logic). The Mcs receive memory requests from the interface units 24, and execute reads and writes as well as providing refresh signals to the DRAMs... to provide a 72 bit data path between the memory array 28 and the interface units 24a, 24b, which utilize an SBC-DBD-SbD ECC scheme, where $b=4$, on a... 26a, 26b to work together and simultaneously supply a 64-bit word to the interface units 24 with minimum latency, one-half of which (DO) comes from the MC 26a, and the other half (D1) comes from the other MC 26b. The interface unit 24 generate and check the ECC check bits. The ECC scheme used will not only... 26 bus 25, as well as in internal registers.

From the viewpoint of the interface units 24, the memory 28 is accessed with two instructions: a "read N doubleword" and a... doubleword read or a block read format. The signal called "data valid" tells the interface units 24 two cycles ahead of time that read data is being returned or not being returned.

As indicated above, the maintenance processor (MP 18; Fig. 1A) has two means of access to the CPUs 12. One is... 18 will write a register contained in the OLAP 285 with instructions that permit the processors 20 to build an image of a sequence of instructions in the memory that will permit them (the processors 20) to commence operation, going to I/O for example to transfer instructions and data from an external (storage) device that will complete the boot process.

The OLAP 285 is also used by the processors 20 to communicate to the MP 18 error indications. For example, if one of the interface units 24 detect a parity error in data received from the memory controller 26, it will... and address transfers on the bus 25 between the MC 26a and the corresponding interface unit 24a. The addressing and data transfers on the DRAM data bus, as well as generation... the CPU 12.

Packet Routing:

The message packets communicated between the various elements of the

processing system 10 (e.g., CPUs 12A, 12B, and devices coupled to the I/O packet... First, each TNet Link L connects to an element (e.g., router 14A) of the processing system 10 via a port that has both receive and transmit capability. Each transmit port... cycle (i.e., each clock period) of the T(underscore)Clk so that the clock synchronization FIFO at the receiving end of the transmission will maintain synchronization.

Clock synchronization is dependent upon the mode in which the processing system 10 is operated. If operating in the simplex mode in which the CPUs 12A... connect directly to the CPUs may drift with respect to each other. Conversely, when the processing system 10 operates in a duplex mode (e.g., the CPUs operate in synchronized, lock-step operation), the clocks between routers 14 and the CPUs 12 to which they... not necessarily phase-locked).

The flow of data packets between the various elements of the processing system 10 is controlled by command symbols, which may appear at any time, even within... ..initiated by a CPU 12, or MP 18, and promulgated to all elements of the processing system 10 by the routers 14 to communicate an event requiring software action by all...command symbol is used in conjunction with near frequency operation as an aid to maintaining synchronization between the two clock signals that (1) transfer each symbol to, and load it in each receiving clock synchronization FIFO, and (2) that retrieves symbols from the FIFO.

SLEEP: This command symbol is sent by any element of the processing system 10 to indicate that no additional packet (after the one currently being transmitted, if... ..received.

SOFT RESET (SRST): The SRST command symbol is used as a trigger during the processes ("synchronization" and "reintegration," described below) that are used to synchronize symbol transfers between the CPUs 12 and the routers 14A, 14B, and then to place... ..SYNC command symbol is sent by a router 14 to the CPU 12 of the processing system 10 (i.e., the sub-processor systems 10A/10B) to establish frequency-lock synchronization between CPUs 12 and routers 14A, 14B prior to entering duplex mode, or when in duplex mode to request synchronization, as will be discussed more fully below. The SYNC command symbol is used in conjunction... ..or duplex to simplex), among other things, as discussed further below in the section on Synchronization and Reintegration.

THIS LINK BAD (TLB): When any system element receiving a symbol from a TNet link L (e.g., a router, a CPU, or an I/O unit) notes an error when receiving a command symbol or packet, it will send a TLB... ..identical pairs of symbols that are compared to one another when pulled from the clock synchronization FIFOs..The DVRG command symbol signals the CPU 12 that a mis-compare has been noted. When received by the CPUs, a divergence detection process is entered whereby a determination is made by the CPUs which CPU may be failing... ..command symbols described above operate to control message flow between the various elements of the processing system 10 (e.g., CPUs 12, router 14, and the like), using principally the BUSY... ..particular TNet port however, an "end node" (i.e., a CPU 12 or I/O unit 17 - Fig. 1) may not assert backpressure because one of its transmit ports is backpressured...Improperly addressed packets are discarded by the router 14.

When a system element of the processing system 10 receives a BUSY command symbol on a TNet link L on which it... ..other command symbols (READY, BUSY, etc.).

Whenever a TNet port of an element of the processing system 10 detects receipt of a READY command symbol, it will terminate transmission of FILL... ..receives.

As will be seen, all elements (e.g., router 14, CPUs 12) of the processing system 10 that connect to a TNet link L for receiving transmitted symbols will receive those symbols via a clock synchronization (CS) FIFO. For example, as discussed above, the interface units 24 of CPUs 12 include all CS FIFOs 102x, 102y (illustrated in Fig. 6). The... ..depth to allow for speed matching, and the elastic FIFOs must provide sufficient depth for processing delays that may occur between transmission of a BUSY command symbol during receipt of a...another data byte in packet B. As packet A progresses to the next router, the process would be repeated. If the router 14 displaces more data bytes than the FIFO can... ..irrespective of its own findings.

SLEEP Protocol:

The SLEEP protocol is initiated by a maintenance processor via a maintenance interface (an on-line access port - OLAP), described below. The SLEEP protocol... ..reintegrate a slice of the system 10. Routers 14 must be idle (no packets in process) in order to change modes without causing data loss or corruption. When a SLEEP command symbol is received, the receiving element of processing system 10 inhibits initiation of transmission of any new packet on the associated transmit port... ..The HALT command symbol provides a mechanism for quickly informing all CPUs 12 in a processing system 10 that is necessary to terminate I/O activity (i.e., message transmissions

between... CPUs that receive HALT command symbols on either of their receive ports (of the interface units 24) will post an interrupt to the interrupt register 280 if the system halt interrupt... interrupt; Fig. 14A).

The CPUs 12 may be provided with the ability to disable HALT processing. Thus, for example, the configuration registers 75 of the interface units 24 can include a "halt enable register" that, when set to a predetermined state (eg., ZERO) disables HALT processing, but reporting detection of a HALT symbol as an error.

Router Architecture:

Referring now to... simplified block diagram of the router 14A is illustrated. The other routers 14 of the processing system 10 (e.g., routers 14B, 14A', etc.) are of substantially identical construction and, therefore... these ports 4, 5 are structured to operate in a frequency locked environment when a processing system 10 is set for duplex mode operation. In addition, when in duplex mode, a... 5025)) will receive the command/data symbols from the CPUs, pass them through the clock synchronization FIFOs 518 (discussed further below), and compare each symbol exiting the clock synchronization FIFOs with a gated compare circuit 517. When duplex operation is entered, a configuration register... 517 to activate the symbol by symbol comparison of the symbols emanating from the two synchronization FIFOs 518 of the router input logic 502 for the ports 4 and 5. Of... to that received, at substantially the same time, by the other port input.

To maintain synchronization in the duplex mode, the two port outputs of the router 14A that transmit to... mode, are duplicated by the routers 14, and returned to both CPUs.) The output logic units 5044), 5045)) that are coupled directly to the CPUs 12 will both receive symbols from... message packet identifies only one of the duplexed CPUs 12, e.g., CPU 12A) in synchronized fashion, presenting those symbols in substantially simultaneous fashion to the two CPUs 12. Of course, the CPUs 12 (more accurately, the associated interface units 24) receive the transmitted symbols with synchronizing FIFOs of substantially the same structure as that illustrated in Fig. 7A so that, even... from the FIFO structures by both CPUs 12 on the same instruction cycle, maintaining the synchronized, lock-step operation of the CPUs 12 required by the duplex operating mode.

As will... conjunction with configuration data written to registers contained in control logic 509 by the maintenance processor 18 (via the on-line access port 285' and serial bus 19A; see Fig. 1A input 502 also assists in maintaining synchronization - at least for those ports sending symbols in the near-frequency environment - by removing received... slower-receiving element receiving symbols from a faster-sending element could overload the input clock synchronization FIFO of the slower-receiving element. That is, if a slower clock is used to pull symbols from the clock synchronization FIFO put there by a faster clock, ultimately the clock synchronization FIFO will overflow.

The preferred technique employed here is to periodically insert SKIP symbols in... stream to avoid, or at least minimize, the possibility of an overflow of the clock synchronization FIFO (i.e., clock synchronization FIFO 518; Fig. 20A) of a router 14 (or CPU 12) due to a T... being slightly higher in frequency than the local clock used to pull symbols from the synchronization FIFO. Using SKIP symbols to by-pass a push (onto the FIFO) operation has the... stall each time a SKIP command symbol is received so that, insofar as the clock synchronization FIFO is concerned, the transmitting clock that accompanied the SKIP symbol was missing.

Thus, logic... the port inputs 502 will recognize, and key off receipt of, SKIP command symbols for synchronization in the near frequency clocking environment so that nothing is pushed onto the FIFO, but... 14, or between routers 14, or between a router 14 and an I/O interface unit 16A - Fig. 1) at a 50 Mhz rate, this allows for a worst case frequency... symbol by supplying FILL or IDLE symbols (which are received and pushed onto the clock synchronization FIFOs, but are not passed to the elastic FIFOs). In short, each elastic FIFO 506... received symbols are then communicated from the input register 516 and applied to a clock synchronization FIFO 518, also by the T(underscore)Clk. The clock synchronization FIFO 518 is logically the same as that illustrated in Figs. 8A and 8B, used in the interface units 24 of the CPUs 12. Here, as Fig. 20A shows, the clock synchronization FIFO 518 comprises a plurality of registers 520 that receive, in parallel, the output of... 516. Associated with each of the

registers 520 is a two-stage validity (V) bit synchronizer 522, shown in greater detail in Fig. 20B, and discussed below. The content of each... registers 520, together with the one-bit content of each associated two-stage validity bit

synchronizer 522, are applied to a multiplexer 524, and the selected register/synchronizer pulled from the FIFO, and coupled to the elastic FIFO 506 by a pair of... is determined the state of the Push Select signal provided by a push pointer logic unit 530; and, selection of which register 520 will supply its content, via the MUX 524... and loading of the register 520 selected by the push pointer logic 530. Similarly, the synchronization FIFO control logic 534 receives the clock signal local to the router (Rec Clk) to... pointer logic 532.

Digressing for a moment, and referring to Fig. 20B, the validity bit synchronizer 522 is shown in greater detail as including a D-type flip-flop 541 with... 530 (Fig. 20A) selects the register 520 of the FIFO with which the validity bit synchronizer is associated for receipt of the next symbol - if not a SKIP symbol.

The delay... Truth Table, below). The D-type flip-flop 543 acts as an additional stage of synchronization, ensuring a stable level at the V output relative to the local Rec Clk. The... flip-flop 542, allowing the Pull signal (a periodic pulse from the sync FIFO Control unit 534) to clear the validity bit on this validity synchronizer 522 when the associated register 520 has been read.

In summary, the validity synchronizer 522 operates... blocked from being routed out a particular port because another message is already in the process of being routed out that port. However, that other message in turn is also blocked packet bound for the CPUs will be replicated by the crossbar logic unit by routing the message packet to both port output 5044)) and 5045)) at the same... P identifies which of path (X or Y) should be used for accessing two sub-processing the device.

The routers 14 provide a capability of constructing a large, versatile routing network for, for example, massively parallel processing architectures. Routers are configured according to their location (i.e., level) in the network by... expansion registers 509j)) and 509k)) are such that bits "def" are used in the algorithmic process, then bits "abc" of the Region ID are compared to the content of the Device... content of the route to default register 509f)) to the final stage of the selection process: check logic 602. Check logic 602 operates to check the status of the port output... a lower level router, and may be located in one or another of the sub-processing systems 10A, 10B. Whether a router is an upper level or lower level router depends... of CPUs 12 and I/O devices 16 to one another, forming a massively parallel processing (MPP) system. Other such MPP systems may exist, and it is those routers configured as... captured. As soon as the message packet's Destination ID is so captured, the selection process begins, proceeding to the development of a target port address that will be used to... an error that will be posted to the MP18 via the router's (or interface unit's) OLAP for action.

Digressing, it should be appreciated that these protocol rules observed by the routers 14 are also observed by the CPUs 12 (i.e., interface units 24) and I/O packet interfaces 17.

Finally, when the router 14A is in the... directly with the CPUs 12A, 12B, and duplex mode is used, a duplex operation logic unit 638 is utilized to coordinate the port output connected to one of the CPUs 12A... was able to write instructions to the OLAP 285 that would be executed by the processors 20 to build a small memory image and routine to permit the CPU 12 to... the clock generation circuit design. There will be one clock generator circuit in each sub-processor system 10A/10B (Fig. 1) to maintain synchronism. Designated generally with the reference numeral 650... used by the various elements (e.g. CPU. 12, routers 14, etc.) of the sub-processor system containing the clock circuit 650 (e.g., 10A).

The clock generator 654 is shown... The 50 Mhz clock signals produced by the counter 663 are distributed throughout the sub-processor system where needed.

Turning now to Fig. 25, there is illustrated the interconnection and use... the clock circuits 650 used to develop synchronous clock signals for a pair of sub-processor systems 10A, 10B (Fig. 1) for frequency locked operation. As illustrated in Fig. 25, the two CPUs 12A and 12B of the sub-processor systems 10A, 10B each have a clock circuit 650, shown in Fig. 25 as clock... 654B of both CPUs 12. A driver and signal line 667 interconnects the two sub-processor systems to deliver the M(underscore)CLK signal developed by the oscillator circuit 652A to the clock generator 654B of the sub-processor system 10B. For fault isolation, and to maintain signal quality, the M(underscore)CLK signal is delivered to the clock generator 654A of the sub-processor system 10A through a separate driver and a loopback connection 668. The reason for the ...the cable (not shown) will establish the connection shown in Fig. 25 between the sub-processor systems 10A, 10B; connected another way, the connections will be similar, but the oscillator 652B... Fig. 25, the M(underscore)CLK signal produced by the oscillator circuit 652A of sub-processing system 10A is used by both sub-processing systems 10A, 10B as their respective SYNC CLK signals and the various other clock signals... produced by the clock generators 654A, 654B. Thereby, the clock signals of the paired sub-processing systems 10A, 10B are synchronized for the frequency locked operation necessary for duplex mode.

The VCXOs 662 of the clock... This allows both clock generators 654A, 654B to continue to provide to the two sub-processing systems 10A, 10B clock signals in the face of improper operation of the oscillator circuit 652A, although the sub-processor systems may no longer be frequency-locked.

The LOCK signals asserted by the phase comparators... LOCK signal signifies that the 50 Mhz signals produced by a clock generator 654 are synchronized, both in phase and in frequency, to the M(underscore)CLK signal. Thus, if either... signal that accompanies the symbol stream, and is used to push symbols onto the clock synchronizing FIFO of the receiving element (router 14, or CPU 12) is substantially identical in frequency... not phase, to that of the receiving element used to pull symbols from the clock synchronization FIFOs. For example, referring to Fig. 23, which illustrates symbols being sent from the router... clock (Local Clk). The former (Rcv Clk) is used to push symbols onto the clock synchronization FIFOs 126 of each CPU, whereas the latter is used to pull symbols from the... much higher frequency clock signal. In such situations provision must be made to ensure that synchronization is maintained between the two CPUs as to symbols pulled from the clock synchronization FIFOs 126 of each.

Here, a constant ratio clocking mechanism is used to control operation of the two clock synchronization FIFOs 126, providing the clock signal that pulls symbols from the two FIFOs at the... control mechanism is shown, designated with the reference numeral 700. As Fig. 26A illustrates, clock synchronization FIFO control mechanism 700 includes an pre-settable, multi-stage serial shift register 702, the... ratio of the clock signal at which symbols are communicated and pushed onto the clock synchronization FIFOs 126 to the frequency of the clock signal used locally. Here, 15 stages are... that will be used as the Local Clk signal to pull symbols from the clock synchronization FIFOs 126, and to operate (update) the pull pointer counter 130. The selected output is... of the CPU 12 to the clock signal used to push symbols onto the clock synchronization FIFO 126, Rcv Clk, the serial shift register is preset so that M stages of... duplexed CPUs 12 with a 50 Mhz clock. Thus, symbols are pushed onto the clock synchronization FIFOs 126 of the CPUs at a 50 Mhz rate. Assume further that the clock... of the MUX 704, which produces the clock signal that pulls symbols from the clock synchronization FIFOs 126, Rcv Clk, will contain, for each 100 ns period, five clock pulses. Thus... five symbols will be pushed onto, and five symbols will be pulled from, the clock synchronization FIFOs 126.

This example is symbolically shown in Fig. 26B, while the timing diagram shown... labelled "IN" in Fig. 27) of the Rcv Clk will push symbols onto the clock synchronization FIFOs 126. During that same 100 ns period, the serial shift register 702 circulates a... clocks which would require additional storage (i.e., an increase in the size of the synchronization FIFO) and impose more latency.

The constant ratio clock circuit presented here (Figs. 26) is... frequency to a clock regime of a different, higher frequency. The use of a clock synchronization FIFO is necessary here for compensating effects of signal delays

when operating in synchronized, duplexed mode to receive pairs of identical command/data symbols from two different sources. However... so long as there are at least two registers in the place of the clock synchronization FIFO. Transferring data from a higher-frequency clock regime to a lower frequency clock regime Packet Interface:

Each of the sub-processor systems 10A, 10B, etc. will have some input/output capability, implemented with various peripheral units, although it is conceivable that the I/O of other sub-processor systems would be available so that a sub-processing system may not necessarily have local I/O. In any event, if local I/O... device (e.g., a signal line) would be received by the I/O packet interface unit 16 and used to form an interrupt packet that is sent to the CPU 12... OLAP bus, configuration information.

On-Line Access Port:

The MP 18 connects to the interface unit 24, memory controller (MC) 26, routers 14, and I/O packet interfaces with interface signals... OLAP 259 is essentially the same, regardless of what element (e.g. router 14, interface unit

24, etc.) it is used with. Fig. 28 diagrammatically illustrates the general structure of the... circuit chip used to implement certain of the elements discussed herein. For example, each interface unit 24, memory controller 26, and router 14 is implemented by an application specific integrated circuit... of the OLAP 259 shown in Fig. 28 describes the OLAP associated with the interface unit 24, the MC 26, and the router 14 of the system.

As Fig. 28 shows... variables, a "soft-vote" (SV) logic element 900 (Fig. 30A) is provided in each interface unit 24 of each CPU 12. As Fig. 30 illustrates, the SV logic elements 900 of each interface unit 24 are connected to one another by a 2-bit SV bus 902, comprising bus lines 902a and 902b. Bus lines 902a carry one-bit values from the interface units 24 of CPU 12A to those of CPU 12B. Conversely, bus line 902b carries one... the CPU 12A.

Illustrated in Fig. 30B, is the SV logic element 900a of interface unit 24a of CPU 12A. Each SV logic element 900 is substantially identical in construction and... 900a should be understood as applying equally to the other logic elements 900a (of interface unit 24b, CPU 12A), and 900b (of the interface units 24a, 24b of CPU 12B) unless noted otherwise. As Fig. 30B illustrates, the SV logic... the logic elements 900a (as well as its own). In this manner the two interface units 24a, 24b of the CPU 12A can communicate asymmetrical variables to each other.

In a... to the remote register 907 of logic element 902a (and that of the other interface unit 24b).

The logic elements 902 form a part of the configuration registers 74 (Fig. 5). Thus, they may be written by the processor unit(s) 20 by communicating the necessary data/address information over at least a portion of... local and remote registers 906 and 907.

The MUX 914 operates to provide each interface unit 24 of CPU 12A with selective use of the bus line 902a for the SV logic elements 900a, or for communicating a BUS ERROR signal if encountered during the reintegration process (described below) used to bring a pair of CPUs 12 into lock-step, duplex operation... same time, write the enable registers 912 of the logic element 900 of both interface units 24 of each CPU. One of the two logic elements 900 of each CPU will... it is the output enable registers 912 associated with the logic elements 900 of interface units 24a of both CPUs 12A, 12B that are written to enable the associated drivers 916. Thus, the output registers 904 of the interface units 24a of each CPU will be communicated to the bus lines 902; that is, the... to the bus line 902a, while the output register associated with logic element 900b, interface unit 24a of CPU 12B is communicated to bus line 902b. The CPUs 12 will both... again written by each CPU, followed again by reading the remote input registers 907. This process is repeated, one bit at a time, until the entire variable is communicated from the... each CPU 12 to the remote input register of the other. Note that both interface units 24 of CPU 12B will receive the bit of asymmetric information.

One example of use... elements 900 are also used to communicate bus errors that may occur during the reintegration process to be described. When reintegration is being conducted, a REINT signal will be asserted. As... ERROR signal is selected by the MUX 914 and communicated to the bus line 902a.

Synchronization:

Proper operation of the sub-processing systems 10A, 10B (Figs. 1A, 2). whether operating independently (simplex mode), or paired and operating in synchronized lock-step (duplex mode), requires assurance that data communicated between the CPUs 12A, 12B and... routers 14A, 14B will be received properly, and that any initial content of the clock synchronization FIFOs 102 (of CPUs 12A, 12B; Fig. 5) and 519 (of routers 14A, 14B; Fig... erroneously interpreted as data or commands. The push and pull pointers of the various clock synchronization FIFOs 102 (in the CPUs 12) and 518 (in the routers 14) need to be... apart, and presetting the associated FIFO queues to some known state. This done, all clock synchronization FIFOs are initialized for near frequency operation. Thus, when the system 10 is initially brought... in order to properly implement the lock-step operation of duplex mode operation, the clock synchronization FIFOs must be synchronized to operate with the particular source from which they receive data in order accommodate any... 14A, 14B to the CPUs 12A, 12B must be accounted for. It is the clock synchronization FIFOs 102 of the paired CPUs ...and present symbols to the two CPUs in a simultaneous manner to maintain lock-step synchronization necessary for duplex mode operation.

In similar fashion, each symbol received by the routers 14A... the CPUs (which is discussed further hereinafter). Again, it is the function of the clock synchronization FIFOs 518 of the routers 14A, 14B that receive message packets from the CPUs 12... so that the symbols received from the two CPUs 12 are retrieved from the clock synchronization FIFOs simultaneously.

Before discussing how the clock synchronization FIFOs of the CPUs and routers are reset, initialized, and synchronized, an understanding of their operation to maintain synchronous lock-step duplex mode operation is believed helpful. Thus, referring for the moment to Fig. 23, the clock synchronization FIFOs 102 of the CPUs 12A, 12B that receive data, for example, from the router... underscore)Clk, from the router 14A to the CPU 12B.

Consider operation of the clock synchronization FIFOs 102x)), 102y)), to receive identical symbol streams during duplex operation. Table 6, below, illustrates... held by the push and pull pointer counters 128, 130 for the CPU 12A (interface unit 24a), and the content of each of the four storage locations (byte 0. ..., byte 3... of Table 6 show the same thing for the FIFO 102y)) of CPU 12B interface unit 24a for each symbol of the duplicated symbol stream.

Assuming the delay 640 is no... " locations of the queues 126. This is because (1) the FIFOs 102 have been synchronized to operate in synchronism (a process described below), and (2) the push pointer counters 128 are clocked by the clock signal... of the symbol stream transmitted by the router 14A will be pulled from the clock synchronization FIFOs 102 of the CPUs 12A, 12B simultaneously, maintaining the required synchronization of received data when operating in duplex mode. In effect, the depths of the queues... order to achieve the operation just described with reference to Table 6, the reset and synchronization process shown in Fig 31A is used. The process not only initializes the clock synchronization FIFOs 102 of the CPUs 12A, 12B for duplex mode operation, but also operates to adjust the clock synchronization FIFOs 518 (Fig. 19A) of the CPU ports of each of the routers 14A, 14B for duplex operation. The reset and synchronization process uses the SYNC command symbol to initiate a time period, delineated by the SYNC CLK signal 970 (Fig. 31B), to reset and initialize the respective clock synchronization FIFOs of the CPUs 12A and 12B and routers 14A, 14B. (The SYNC CLK signal... It is of a lower frequency than that used to receive symbols by the clock synchronization FIFOs, T(underscore)Clk. For example, where T(underscore)Clk is approximately 50 MHz, the... signal is approximately 3.125 MHz.)

Turning now to Fig. 31A, the reset and initialization process begins at step 950 by switching the clock signals used by the CPUs 12A, 12B and routers 14A, 14B as the transmit (T(underscore)Clk) and the unit's local clock (Local Clk) clock signals so that they are derived from the same... In addition, configuration registers in the CPUs 12A,

12B (configuration registers 74 in the interface units 24) and the routers 14A, 14B (contained in control logic unit 509 of routers 14A, 14B) are set to the FreqLock state.

The following discussion involves step 952, and makes reference to the interface unit 24 (Fig.5), router 14A (Fig. 19A) and Figs. 31A and 31B. With the clock... otherwise be sent followed by a self-addressed message packet. Any message packet in the process of being received and retransmitted when the SLEEP command symbols are received and recognized by... per the destination address). The SLEEP command symbol operates to "quieece" router 14A for the synchronization process. The self-addressed message packet sent by the CPU 12A, when received back by the... message packet sent after the SLEEP command symbol would necessarily have to be the last processed by the router 14A.

At step 954 the CPU 12A checks to see if it... the router will assert a RESET signal 972 that is applied to the two clock synchronization FIFOs 518 contained in the input logic 5054)), 5055)) of the router that receive symbols directly from CPUs 12A, 12B. RESET, while asserted, will hold the two clock synchronization FIFOs 518 in a temporarily non-operating reset state with the push and pull pointer... As each of the CPUs 12 receive SYNC symbols are detected by the storage and processing units of the packet receivers 96 (Figs. 5 and 6) cause the RESET signal to be asserted by the packet receivers 96 (actually, storage and processing elements 110; Fig. 6) of each CPU 12. the RESET signal is applied to the...t4))), CPUs 12 and routers 14A, 14B de-assert the RESET signals, and the clock synchronization FIFOs of the CPUs 12A, 12, and routers 14A, 14B are released from their reset... the delay, the router 14A and CPUs 12 resume pulling data from their respective clock synchronization FIFOs and resume normal operation. The clock synchronization FIFOs of the router 14A begin pulling symbols from the queue (previously set by RESET... from the CPU 12A with the T(underscore)Clk will be pushed onto the clock synchronization FIFO at, for example, queue location 0 (or whatever other location pointed to by the... 0 (or whatever other location the push pointer was set to by RESET). The clock

synchronization FIFOs of the router 14A are now synchronized to accommodate whatever delay 640 may be present in one communications path, relative to the... and the CPUs 12A, 12B.

Similarly, at the same virtual time, operation of the clock synchronization FIFOs 102 of both CPUs 12A, 12B is resumed, synchronizing them to the router 14A. Also, the CPUs 12A, 12B quit sending the SLEEP command... in favor of READY symbols, and resume message packet transmission, as appropriate.

That completes the synchronization process for the router 14A. However, the process must also be performed for the router 14B. Thus, the CPU 12A returns to step... however, assuming that the CPUs 12A, 12B are operating in duplex mode, the method and apparatus used to detect and handle a possible error, resulting in divergence of the CPUs from... via a message packet destined for a peripheral device of one or the other sub-processor systems 10A, 10B. Depending upon the destination of the outgoing message packet, step 1002 will... router 14 will issue an ERROR signal to the router control logic 509, causing the process to move to step 1004 where the router 14 detecting divergence will transmit a DVRG...time outs to occur. A router detecting divergence (without also detecting any simple link error) buys itself time to check the CRC of the received message packet by waiting for the... router 14, or received, all further message packets received from the CPUs and in the process of being routed when divergence was detected, or the DVRG symbol received, will be passed...1010) contained in a one of the configuration registers 74 (Fig. 5) of the interface unit 24 of each CPU.

Returning for the moment to step 1006, the determination of which... local" is meant to refer to the router 14A, 14B contained in the same sub-processor system 10A, 10B as the CPU. For example, referring to Fig. 1A, router 14A is... bit mentioned above: the bit contained in one of the configuration registers 74 of interface unit 24(Fig. 5) of each CPU. When set to a first state, that particular CPU...the other CPU. In response, the state machines (not shown) within the control and status unit 509 (Fig. 19A) changes the "favorite" bits described above.

A few examples may facilitate understanding... ..DVRG symbol will echo that symbol to the routers 14A, 14B, start its internal divergence process timer, and begin determination of whether to continue or terminate. Having received a TLB symbol...to diverge with no errors reported. This can happen only if software (running on the processors 20) uses known divergent data to alter state. For example, suppose each CPU 12 has... ..number of the CPU 12A will differ from that of the CPU 12B. If the processors use the serial number to change the sequence of instructions executed (say, by branching if the serial number comes after some value) or to modify the value contained in a processor register, the complete "state" of the CPUs 12 will differ. In such cases, the "asymmetrical... ..of the primary CPU simply allows one CPU, and thereby the system 10, to continue processing without software intervention.

- An error at the output of the interface unit 24 of a CPU 12 will be detected by the router 14A, 14B, depending upon... ..router 14A, 14B that connects to a CPU 12 will be detected by the interface unit 24 of the affected CPU. The CPU will send a TLB symbol to the faulty... ..possible failure and, without external intervention, and transparently to the system user, remove the failing unit (CPU 12A or 12B, or router 14A or 14B) from the system to obviate or... ..reintegration." The discussion will refer to the CPUs 12A, 12B, routers 14A, 14B, and maintenance processor 18A, 18B shown forming parts of the processing system 10 illustrated in Fig. 1A. In addition, discussion will refer to the processors 20a, 20b, the interface units 24a, 24b, and the memory controllers 26a, 26b (Fig. 2) of the CPUs 12A, 12B as single units, since that is the way they function.

Reintegration is used to place two CPUs in... ..both of the paired CPUs at virtually the same time.

The major steps in the process for changing from simplex mode operation of the one on-line CPU to duplex mode... ..greater detail by the flow diagrams of Figs. 33A - 33D, generally are:

1. Setup and synchronize the two CPUs (one on-line, the other off-line) and their connected routers to... ..the memory of the on-line CPU to the off-line CPU, maintaining a tracking process that monitors changes in the memory of the on-line CPU that have not been... ..and may need to be copied over to, the off-line CPU;
3. Setup and synchronize the CPUs to run a delayed (slave) duplex mode from the same instruction stream (lock...will write the predetermined registers (not shown) of the control registers 74 in the interface units 24 of CPUs 12A and 12B, to a next state (after a soft operation) in... ..the off-line CPU 12B.

Next, a sequence is entered (steps 1060 - 1070) that will synchronize the clock synchronization FIFOs of the CPUs 12A, 12B and routers 14A, 14B in much the same fashion... ..the same steps described above in connection with the discussion of Figs. 31A, 31B to synchronize the clock synchronization FIFOs. The on-line CPU 12A will send the sequence of a SLEEP symbol, self-addressed message packet, and SYNC symbol which, with the SYNC CLK signal, operates to synchronize CPUs and routers. Once so synchronized, the on-line CPU 12A then, at step 1066, sends a Soft Reset (SRST) command... ..of all configuration registers and control registers (e.g., configuration registers 74 of the interface units 24) cache, and the like to memory 28 of the on-line CPU, copying the... ..time to have the system 10 off-line for reintegration. For that reason, the reintegration process is performed in a manner that allows the on-line CPU to continue executing user... ..not match that of the off-line CPU. The reason for this is that normal processing by the processor 20 of the on-line CPU can change memory content after it has been copied... ..when a memory location is written in the on-line CPU 12A during the reintegration process it is marked as "dirty;" second, all copying of memory to the off-line CPU...may, however, limit the ability to detect two-bit errors. But, since the memory copying process will last for a only relatively short period of time, this risk is believed acceptable... ..memory location in CPU 12A is made (either an incoming I/O write, or a processor write operation). The returning data (that was copied over to the off-line CPU) would... ..controller 26 (Fig. 2) of the on-line CPU to monitor memory locations in the process of being copied over to the off-line CPU 12B. The memory controller uses a...within the block had been written by another operation (e.g., a write by the processor 20, an I/O write, etc.), that prior write operation will flag the location in... ..still must be copied over to the off-line CPU 12B.

Returning to the reintegration process, and now to Fig. 33B, the memory tracking (AtomicWrite mechanism and using ECC to mark... ..entails writing a reintegration register (not shown; one of the configuration registers 74 of interface unit 24 - Fig. 5) to cause a reintegration (REINT) signal to be asserted. The REINT signal... ..is left alone. Throughout the incremental copy operations, the normal actions of the on-line processor will mark some memory locations dirty.

Several passes of incremental copying will need to be... ..the number of successful WriteConditional operations at the end of each pass through memory, the processors 20 can determine the effect of a given pass compared to the previous pass. When the benefits drop off, the processors 20 will give up on the precopy operations. At this point the reintegration process is ready to place the two CPUs 12A, 12B into lock-step operation.

Thus, the... ..in Fig. 33C, where at step 1100, the on-line CPU 12A momentarily halts foreground processing, i.e., execution of a user application. The remaining state (e.g., configuration registers, cache, etc.) of the on-line processors 20 and its caches is then read and written to a buffer (series of memory... ..to the off-line CPU 12B, together with a "reset vector" that will direct the processor units 20 of both CPUs 12A, 12B to a reset instruction.

Next, step 1106 will quiesce... ..to ensure that the FIFOs of the routers are clear, that the FIFOs of the processor interfaces 24 are clear, and no further incoming I/O message packets are forthcoming. At... ..symbol will be received and acted upon by both CPUs 12A, 12B, to cause the processor units 20 of each CPU to jump to the location in memory 28 containing the reset... ..a subroutine that will restore the stored state of both CPUs 12A, 12B to the processor units 20, caches 22, registers, etc. The CPUs 12A, 12B will then begin executing the same...enabling of the ECC bit to mark dirty locations must now be disabled, since the processors are doing the same thing to the same memory. During this stage of the reintegration... ..encountered by CPU 12A.

Meanwhile, the bus error in the CPU 12A will cause the processor unit 20 to be forced into an error-handling routine to determine (1) the cause of... ..error was caused by an attempt to read a memory location marked dirty. Accordingly, the processor unit 20 will initiate (via the BTE 88 -- Fig. 5) the AtomicWrite mechanism to copy the... ..the SRST symbols are now received by the CPUs 12A, 12B, they will cause both processor units 20 of the CPUs to be reset to start from the same location with the... ..will periodically update, e.g., a database or audit file that is indicative of the processing of the primary CPU up to that point in time of the update. Should the...in error-checking redundancy to the CPU 12B, in the same manner that the individual processor units 20a, 20b of the CPU 12A provide fail-fast, fault tolerance for the CPU - when... ..cost system is applicable, as illustrated in Fig. 34. As shown in Fig. 34, a

processing system 10' includes the CPU 12A and routers 14A, 14B structured as described above. The... ..and the CPUs are also the same.

Thus, the CPU 12B' comprises only a single processor unit 20' and associated support components, including the cache 22', interface unit (IU) 24', memory controller 26', and memory 28'. Thus, while the CPU 12A is structured in the manner shown in Fig. 2, with cache processor unit, interface unit, and memory control redundancies, approximately one-half of those components are needed to implement CPU... ..stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... ..i.e., parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of... ..inputs connected to receive the data output from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... ..for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the... ..attached, without the assistance of any other CPU 12 in the system. Many prior parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific

processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O device becomes... ..rest of the system. Other prior systems provide access to I/O through pairs of processors so that should one of the processors fail, access to the corresponding I/O is still available through the remaining I/O... ..if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... ..multiprocessing system access to every peripheral , as done here, operates to extend the "primary"/"backup" process taught in the above-identified U.S. Patent No. 4,228,496. There, a multiple CPU system may have a primary process running on one CPU, while a backup process resides in the background on another of the CPUs. Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be... ..e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every ...CPU can access any I/O without the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... ..four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... ..an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time...indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... ..When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... ..response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Claims: ...B1

1. A fault tolerant processing system (10), comprising:

a first central processing unit (12A) comprising a pair of first processor devices (20a, 20b) operating to execute each instruction of an instruction stream at substantially the same moment in time;

a second central processing unit (12B) comprising a second processor device (20a, 20b) operating to execute each instruction of a substantially identical copy of the instruction stream, the pair of first processor devices and the second processor device executing identical instructions of the instruction stream and the identical copy of the instruction stream at substantially the same moment in time, whereby the first and second central processing units operate in synchronism to perform substantially the same operations at substantially the same moments in... ..14b) connected to receive and compare the output data from the first and second central processing units, and including means for issuing a divergence signal to each central processing unit when a mis-compare in said output data is detected; and

means, in each of said first and second central processing units for detecting said divergence signal and based on a number of possible tabulated conditions that may be detected by the central processing units, determining if said respective central processing unit should terminate operation in response to said divergence signal.

2. The fault tolerant processing system of claim 1, wherein the data checking element is a data communicating element (14a... ..and first and second inputs (Lx, Ly) respectively connected to the first and second central processing units (12A, 12B), the data communicating element operating to receive and forward output data packets from the first and second central processing units and, following said determining which processing unit should terminate operation, terminating said forwarded output data packets with either a packet good status indicator (TPG) or packet bad status indicator (TPB).

3. The fault tolerant processing system of claim 2 in which the data communicating element is adapted to continue forwarding... ..status indicator (TPB) depending upon the outcome of said determining operation.

4. The fault tolerant processing system of claim 1, claim 2 or claim 3 including a data sending device (16... ..received input data and transmits the received input data to the first and second central processing units at substantially the same time.

5. The fault tolerant processing system of claim 1 in which each central processing unit is coupled to a local one and a remote one of said data checking elements and includes a said means for detecting said divergence signals therefrom.

6. The fault tolerant processing system of claim 5 in which each said means for detecting said divergence signals determines whether to terminate its respective processing unit according to link status indicators (TLB, OLB) supplied by each of said local one and remote one of said data checking elements.

7. The fault tolerant processing system of claim 1 wherein said means for detecting said divergence signal and determining which processing unit in said processing system should terminate operation comprises means for switching said fault tolerant processing system from a duplex mode of operation in which said first and second processing units are executing each instruction of an instruction stream at substantially the same time to a simplex mode of operation in which only said non-terminated processing unit continues processing said instruction stream.

Claims: ...a sensiblement le meme instant ;

une seconde unite centrale de traitement (12B) comprenant un second processeur (20a, 20b) fonctionnant pour executer chaque instruction d'une copie sensiblement identique du train d'instructions, la paire de premiers processeurs et le second processeur executant des instructions identiques du train d'instructions et de la copie identique du train...

6/K/20 (Item 20 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Patent Assignee:

- Compaq Computer Corporation... ;
;;

Country	Number	Kind	Date	
Type		Pub. Date	Kind	Text
...Compaq Computer Corporation (687797) 20555 SH 249 Houston, Texas 77070-2698 US		19		
Examination...		19		
Available Text		Language	Update	Word Count
Total Word Count (Document A)				
Total Word Count (Document B)				
Total Word Count (All Documents)				

Specification: ...and 08/485,055 filed concurrently herewith.

The present invention is directed generally to data processing systems, and more particularly to a multiple processing system and a reliable system area network that provides connectivity for interprocessor and input/output... ..and communications systems to general purpose high availability commercial systems. The evolution of fault tolerant computers has been well documented (see D. P. Siewiorek, R. S. Swarz, "The Theory and Practice... ..and the Jet Propulsion laboratory began to apply fault tolerance to the development of guidance computers for aerospace applications. The 1960's also saw the development of the first AT&T electronic switching systems.

The first commercial fault tolerant machines were introduced by Tandem Computers in the 1970's for use in on-line transaction processing applications (J. Bartlett, "A NonStop Kernal," in proc. Eighth Symposium on Operating System Principles, pp... ..systems were introduced in the 1980's (O. Serlin, "Fault- Tolerant Systems in Commercial Applications," Computer, pp. 19-30, August 1984). Current commercial fault tolerant systems include distributed memory multi-processors, shared-memory transaction based systems, "pair-and-spare" hardware fault tolerant systems (see R. Freiburghouse, "Making Processing Fail-safe," Mini-micro Systems, pp. 255-264, May 1982; U.S. Patent No. 4... ..system.), and triple-modular-redundant systems such as the "Integrity" computing system manufactured by Tandem Computers Incorporated of Cupertino, California, assignee of this application and the invention disclosed herein.

Most applications of commercial fault tolerant computers fall into the category of on-line transaction processing. Financial institutions require high availability for electronic funds transfer, control of automatic teller machines, and... ..telecommunications systems.

Vendors of fault tolerant machines attempt to achieve both increased system availability, continuous processing, and correctness of data even in the presence of faults. Depending upon the particular system architecture, application software ("processes") running on the system either continue to run despite failures, or the processes are automatically restarted from a recent checkpoint when a fault is encountered. Some fault tolerant systems are provided with sufficient component redundancy to be able reconfigure around failed components, but processes running in the failed modules are lost. Vendors of commercial fault tolerant systems have extended fault tolerance beyond the processors and disks. To make large improvements in reliability, all sources of failure must be addressed... ..power supplies, fans and inter-module connections.

The "NonStop," and "Integrity" architectures manufactured by Tandem Computers Incorporated, (both respectively illustrated broadly in U.S. Patent No. 4,228,496 and U... ..assigned to the assignee of this application; NonStop and Integrity are registered trademarks of Tandem Computers Incorporated) represent two current approaches to commercial fault tolerant computing. The NonStop system, as generally... ..above-identified U.S. Patent No. 4,278,496, employs an architecture that uses multiple processor systems designed to continue operation despite the failure of any single hardware component. In normal operation, each processor system uses its major components independently and concurrently, rather than as "hot backups". The NonStop system architecture may consist of up to

16 processor systems interconnected by a bus for interprocessor communication. Each processor system has its own memory which contains a copy of a message-based operating system. Each processor system controls one or more input/output (I/O) busses. Dual-porting of I/O controllers and devices provides multiple paths to each device. External storage (to the processor system), such as disk storage, may be mirrored to maintain redundant permanent data storage.

This... hardware, while fault recovery is the responsibility of the software.

Also, in the Nonstop multi-processor architecture, application software ("process") may run on the system under the operating system as "process-pairs," including a primary process and a backup process. The primary process runs on one of the multiple processors while the backup process runs on a different processor. The backup process is usually dormant, but periodically updates its state in response to checkpoint messages from the primary process. The content of a checkpoint message can take the form of complete state update, or currently most application code runs under transaction processing software which provides recovery through a combination of checkpoints and transaction two-phase commit protocols.

Interprocessor message traffic in the Tandem Nonstop architecture includes each processor periodically broadcasting an "I'm Alive" message for receipt by all the processors of the system, including itself, informing the other processors that the broadcasting processor is still functioning. When a processor fails, that failure will be announced and identified by the absence of the failed processor's periodic "I'm Alive" message. In response, the operating system will direct the appropriate backup processes to begin primary execution from the last checkpoint. New backup processes may be started in another processor, or the process may be run with no backup until the hardware has been repaired. U.S. Patent... example of this technique.

Each I/O controller is managed by one of the two processors to which it is attached. Management of the controller is periodically switched between the processors. If the managing processor fails, ownership of the controller is automatically switched to the other processor. If the controller fails, access to the data is maintained through another controller.

In addition to providing hardware fault tolerance, the processor pairs of the above-described architecture provide some measure of software fault tolerance. When a processor fails due to a software error, the backup processor frequently is able to successfully continue processing without encountering the same error. The software environment in the backup processor typically has different queue lengths, table sizes, and process mixes. Since most of the software bugs escaping the software quality assurance tests involve infrequent data dependent boundary conditions, the backup processes often succeed.

In contrast to the above-described architecture, the Integrity system illustrates another approach... fault recovery is the logical choice since few modifications to the software are required. The processors and local memories are configured using triple-modular-redundancy (TMR). All processors run the same code stream, but clocking of each module is independent to provide tolerance... three streams is asynchronous, and may drift several clock periods apart. The streams are re-synchronized periodically and during access of global memory. Voters on the TMR Controller boards detect and mask failures in a processor module. Memory is partitioned between the local memory on the triplicated processor boards and the global memory on the duplicated TMRC boards. The duplicated portions of the... techniques to detect failures. Each global memory is dual ported and is interfaced to the processors as well to the I/O Processors (IOPs). Standard VME peripheral controllers are interfaced to a pair of busses through a Bus... the BIMs to switch control of all controllers to the remaining IOP. Mirrored disk storage units may be attached to two different VME controllers.

In the Integrity system all hardware failures... reintegrated on-line.

The preceding examples illustrate present approaches to incorporating fault tolerance into data processing systems. Approaches involving software recovery require less redundant hardware, and offer the potential for some... have been developed on other systems.

Thus, the systems described above provide fault tolerant data processing either by hardware (e.g, fail-functional, employing redundancy) or by software techniques (fail-fast... hardware). However, none of the systems described are believed capable of providing fault tolerant data processing, using both hardware (fail-functional) and software (fail-fast) approaches, by a single data processing system.

Computing systems, such as those described above, are often used for electronic commerce: electronic data interchange (EDI) and global messaging. Today's demands upon such electronic commerce, however, is demanding more and more throughput capacity as the number of users increases and... networks such as local area networks (LANS), and the like.

A key requirement for a server architecture is the ability to move massive quantities of data. The server should have high bandwidth that is scalable, so that added throughput capacity can be added... response time, latency affects service levels and employee productivity.

The present invention provides a multiple-processor system that combines both of the two above-described approaches to fault tolerant architecture, hardware redundancy and software recovery techniques, in a single system.

Broadly, the present invention includes a processing system composed of multiple sub-processing systems. Each sub-processing system has, as the main processing element, a central processing unit (CPU) that in turn comprises a pair of processors operating in lock-step, synchronized fashion ...execute each instruction of an instruction stream at the same time. Each of the sub-processing systems further include an input/output (I/O) system area network system that provides redundant communication paths between various components of the larger processing

system, including a CPU and assorted peripheral devices (e.g., mass storage

units, printers, and the like) of a sub-processing system, as well as between the sub-processors that may make up the larger overall processing system. Communication between any component of the processing system (e.g., a CPU and another CPU, or a CPU and any peripheral device, regardless of which sub-processing system it may belong to) is implemented by forming and transmitting packetized messages that are... responsible for choosing the proper or available communication paths from a transmitting component of the processing system to a destination component based upon information contained in the message packet. Thus, the... peripherals, but permits it to also be used for interprocessor communications.

As indicated above, the processing system of the present invention is structured to provide fault-tolerant operation through both "fail... at a variety of points in the various data paths between the (lock-step operated) processor elements of the CPU and its associated memory. In particular, the processing system of the present invention conducts error-checking at an interface, and in a manner... little impact on performance. Prior art systems typically implement error-checking by running pairs of processors, and checking (comparing) the data and instruction flow between the processors and a cache memory. This technique of error-checking tended to add delay to the... error-checking precluded use of off-the-shelf parts that may be available (i.e., processor/cache memory combinations on a single semiconductor chip or module). The present invention performs error-checking of the processors at points that operate at slower rates, such as the main memory and I/O interfaces which operate at slower speeds than the processor-cache interface. In addition, the error-checking is performed at locations that allow detection of errors that may occur in the processors, their cache memory, and the I/O and memory interfaces. This allows simpler designs for... other data integrity checks.

Error-checking of the communication flow between the components of the processing system is achieved by adding a cyclic-redundancy-check (CRC) to the message packets that... "Good" (TPG) or "This Packet Bad" (TPB) - is appended to every packet. A maintenance diagnostic processor can use this information to isolate a link or router element that introduces an error... of topologies, so that alternate paths can be provided between any two elements of a processing system (e.g., between a CPU and an I/O device), for communication in the... so (e.g., by creating a "deadlock" condition, discussed further below).

The CPUs of a processing system are capable of operating in one of two basic modes: a "simplex mode" in... independently of the other, or a "duplex" mode in which pairs of CPUs operate in synchronized, lock-step fashion. Simplex mode operation provides the capability of recovering from faults that are... U.S. Pat. No. 4,228,496 which teaches a multiprocessing system in which each processor has the capability of checking on the operability of its sibling processors, and of taking over the processing of a processor found or believed to have failed). When operating in duplex mode, the paired CPUs both... fault tolerant platform for less robust operating systems (e.g., the UNIX operating system). The processing system of the present invention, with the paired, lock-step CPUs, is structured so that... masked (i.e., operating despite the existence of a fault), primarily through hardware.

When the processing system is operating in duplex mode, each CPU pair uses the I/O system to access any peripheral of the processing system, regardless of which (of the two, or more) sub-processor system the peripheral may be ostensibly a member of. Also, in duplex mode, message packets... message for the CPU pair (from either a peripheral device such as a mass storage unit or from a processing unit), will replicate the message and deliver it to both CPUs of the pair using synchronization methods that ensure that the CPUs remain synchronized. In effect, the duplex CPU pair, as viewed from the I/O system and other... as a single CPU. Thus, the I/O system, which includes elements from all sub-processing systems, is made to be seen by the duplex CPU pair as one homogeneous system... a multiprocessor system in which the CPU of any one is actually a pair of synchronized, lock-step CPUs.

Yet another important aspect of the present invention is that interrupts issuing... interrupts via the message packet system ensures that they will arrive at duplexed CPUs in synchronized fashion, in the same manner as I/O message packets. Interrupt message packets will contain... the system. In addition, using the same messaging system to communicate data between I/O units and the CPUs and to communicate interrupts to the CPUs preserves the ordering of I... the implementation of a technique of validating access to the memory of any CPU. The processing system, as structured according to the present invention, permits the memory of any CPU to a CPU and any other component of the processor system. Thereby, the individual processor units of the CPU are removed from the more mundane tasks of getting information from memory and out onto the TNet network, or accepting information from the network. The processor unit of the CPU merely sets up data structures in memory containing the data to be... is required, where in memory the response is to be placed when received. When the processor unit completes the task of creating the data structure, the block transfer engine is notified to... response is received, it is routed to the expected memory location identified, and notifies the processor unit that the response was received.

Further aspects and features of the present invention will become... invention, which should be taken in conjunction with the accompanying drawings.

Fig. 1A illustrates a processing system constructed in accordance with the teachings of the present invention, and Figs. 1B and 1C illustrate two alternate configurations of the processing system of Fig. 1A, employing clusters or arrangements of the processing system of Fig. 1A;

Fig. 2 illustrates, in simplified block diagram form, the central processing unit (CPU) that forms a part of each sub-processor system of Figs. 1A - 1C;

Figs. 3A - 3D and 4A - 4C illustrate the construction of... the area network I/O system shown in Fig. 2;

Fig. 5 illustrates the interface unit that forms a part of the CPUs of Fig. 2 to interface the processor and memory with the I/O area network system;

Fig. 6 is a block diagram, illustrating a portion of packet receiver of the interface unit of Fig. 5;

Fig. 7A diagrammatically illustrates the clock synchronization FIFO (CS FIFO) used by the packet receiver section packet receiver shown in Fig. 6;

Fig. 7B is an block diagram of a construction of the clock synchronization FIFO structure shown in Fig. 7A;

Fig. 8 illustrates the cross-connections for error-checking outbound transmissions from the two interface units of a CPU;

Fig. 9 illustrates an encoded (8B to 9B) data/command symbol;

Fig. 10 illustrates the method and structure used by the interface unit of Fig. 5 to cross-check for errors data being transferred to the memory controllers... ..of a CPU of Fig. 2 to other (external to the CPU) components of the processing system;

Fig. 12 is a block diagram that diagrammatically illustrates the formation of an address... ..14A illustrates the logic for posting interrupt requests to queues in memory and to the processor units of the CPU of Fig. 2;

Fig. 14B illustrates the process used to form a memory address for a queue entry;

Fig. 15 is a block... ..data output constructs formed in the memory of the CPU of Fig. 2 by a processor unit, and containing data to be sent via the area I/O networks shown in Figs. 1A - 1C, and also illustrating the block transfer engine (BTE) unit of the interface unit of Fig. 5 that operates to access the data output constructs for transmission to the... ..pair of memory controllers between memory of a CPU of Fig. 2 and its interface unit for accessing from memory 72 bits of data, including two simultaneously-accessed 32-bit words... ..other for error-checking;

Fig. 19A is a simplified block diagram illustration of the router unit used in the area input/output networks of the processing systems shown in Figs. 1A - 1C;

Fig. 19B illustrates comparison on two port inputs of the router unit of Fig. 19A;

Fig. 20A is a block diagram the construction of one of the six input ports of the router unit shown in Fig. 19A;

Fig. 20B is a block diagram of the synchronization logic used to validate command/data symbols received at an input port of the router unit of Fig. 19A;

Fig. 21A is a block diagram illustration of the target port selection... ..is a block diagram illustration of one of the six output ports of the router unit shown in Fig. 19A;

Fig. 23 is an illustration of the method used to transmit identical information to a duplexed pair CPUs of Fig. 2 in synchronized fashion when the processing system is operating in lock-step (duplex) mode, using a pair the FIFOs of Fig... ..is a simplified block diagram illustrating the clock generation system of each of the sub-processing systems of Figs. 1A - 1C for developing the plurality of clock signals used to operate the various elements of that sub-processing system;

Fig. 25 illustrates the topology used to interconnect the clock generation systems of paired sub-processing systems for synchronizing the various clock signals of the pair of sub-processing systems to one another;

Fig. 26A and 26B illustrates a FIFO constant rate clock control logic used to control the clock synchronization FIFO of Figs. 8 or 20 in the situation when the two clocks used to... ..structure of the on-line access port (OLAP) used to provide access to the maintenance

processor (MP) to the various elements of the system of Fig. 1A (or those of Figs... ..the soft-flag logic used to handle asymmetric variables between the CPUs of paired sub-processing systems operating in duplex mode;

Fig. 31A shows a flow diagram, and Fig. 31B illustrates a portion of SYNC CLK, both of which are used to reset and synchronize the clock synchronization FIFOs of the CPUs and routers of the processing system of Fig. 1A that receive information from each other;

Fig. 32 is a flow... 33A - 33D generally illustrate the procedure used to bring an one of the CPUs of processing system shown in Fig. 1A into lock-step, duplex mode operation with the other of the CPUs without measurably halting operation of the processing system; and

Fig. 34 illustrates a reduced cost architecture incorporating teachings of the invention; and... to the figures and, for the moment, principally Fig. 1A, there is illustrated a data processing system, designated with the reference 10, constructed according to the various teachings of the present invention. As Fig. 1A shows, the data processing system 10 comprises two sub-processor systems 10A and 10B each of which are substantially the same in structure and function... should be appreciated that, unless noted otherwise, a description of any one of the sub-processor systems 10 will apply equally to any other sub-processor system 10.

Continuing with Fig. 1A therefore, each of the sub-processor systems 10A, 10B is illustrated as including a central processing unit (CPU) 12, a router 14, and a plurality of input/output (I/O) packet interfaces... one of the I/O packet interfaces 16 will also have coupled thereto a maintenance processor (MP) 18.

The MP 18 of each sub-processor system 10A, 10B connects to each of the elements of that sub-processor system via an IEEE 1149.1 test bus 17 (shown in phantom in Fig. 1A... accompanying clock signal. As Fig. 1A further illustrates, TNet Links L also interconnect the sub-processor systems 10A and 10B to one another, providing each sub-processor system 10 with access to the I/O devices of the other as well as inter-CPU communication. As will be seen, any CPU 12 of the processing system 10 can be given access to the memory of any other CPU 12, although... the memory of a CPU 12 by a wayward peripheral device 17.

Preferably, the sub-processor systems 10A/10B are paired as illustrated in Fig. 1A (and Figs 1B and 1C, discussed below), and each sub-processor system 10A/10B pair (i.e., comprising a CPU 12, at least one router 14... 12A) connects, by a TNet Link L to a router (14A) of the corresponding sub-processor system (e.g., 10A). Conversely, the Y port connects the CPU (12A) to the router (14B) of the companion sub-processor system (10B). This latter connection not only provides a communication path for access by a CPU (12A) to the I/O devices of the other sub-processor system (10B), but also to the CPU (12B) of that system for inter-CPU communication.

Information is communicated between any element of the processing system 10 and any other element (e.g., CPU 12A of sub-processor system 10A) of the system and any other element of the system (e.g., an I/O device associated with an I/O packet interface 16B of sub-processor system 10B) via message "packets." Each message packet is made up of a number of... this reason, a unique method of receiving the symbols at the receiver, using a clock synchronization first-in-first-out (CS FIFO) storage structure (described more fully below), has been developed... operation means just that: the frequencies of the clock signals of the transmitter and receiver units are locked, although not necessarily in phase. Frequency locked clock signals are used to transmit symbols between the routers 14A, 14B and the CPUs 12 of paired sub-processor systems (e.g., sub-processor systems 10A, 10B, Fig. 1A). Since the clocks of the transmitting and receiving element are not phase related, a clock synchronization FIFO is again used -- albeit operating in a slightly different mode from that used for... difference, as will be seen, is due to the fact that pairs of the sub-processor systems 10 can be operated in a synchronized, lock-step mode, called duplex mode, in which each CPU 12 operates to execute the... 1A illustrates another feature of the invention: a cross-link connection between the two sub-processor systems 10A, 10B through the use of additional routers 14 (identified in Fig. 1A as... RY(sub(1)), and RY(sub(2)) form a cross-link connection between the sub-processors 10A, 10B (or, as shown, "sides" X and Y, respectively) to couple them to L... the routers RX(sub(2)) and RY(sub(2)) provide the I/O packet interface units 16x and 16y with a dual ported interface. Of course, it will now be evident... lend themselves to being used in a manner that can extend the configuration of the processing system 10 to include additional sub-processor systems such as illustrated in Figs. 1B and 1C. In Fig. 1B, for example, one... of each of

the routers 14A and 14B is used to connect the corresponding sub-processor systems 10A and 10B to additional sub-processor systems 10A' and 10B' forming thereby a larger processing system comprising clusters of the basic processing system 10 of Fig. 1.

Similarly, in Fig. 1C the above concept is extended to form an eight sub-processor system cluster, comprising sub-processor systems pairs 10A/10B, 10A'/10B', 10A''/10B'', and 10A'''/10B'''. In turn, each of the sub-processor systems (e.g., sub-processor system 10A) will have essentially the same basic minimum configuration of a CPU 12, a... ..by a I/O packet interface 16, except that, as Fig. 1C shows, the sub-processor systems 10A and 10B include additional routers 14C and 14D, respectively, in order to extend the cluster beyond sub-processor systems 10A'/10B' to the sub-processor systems 10A''/10B'' and 10A'''/10B''''. As Fig. 1C further illustrates, unused ports 4 and... ..the routers 14 when configuring the topology of the system 10, any CPU 12 of processing system 10 of Fig. 1C can access any other "end unit" (e.g., a CPU or I/O device) of any of the other sub-processor systems. Two paths are available from any CPU 12 to the last router 14 connecting to the I/O packet interface 16. For example, the CPU 12B of the sub-processor system 10B' can access the I/O 16''' of sub-processor system 10A''' via router 14B (of sub-processor system 10B'), router 14D, and router 14B (of sub-system 10B''') and, via link LA... ..10A'''), OR via router 14A (of sub-system 10A'), router 14C, and router 14A (sub-processor system 10A'''). Similarly, CPU 12A of sub-processor system 10A'' may access (via two paths) memory contained in the CPU 12B of sub-processor 10B to read or write data. (Memory accesses by one CPU 12 of another component of the processing system requires, as will be seen, the components seeking access to have authorization to do... ..prevents corruption of memory data of a CPU by erroneous access.)

The topology of the processing system shown in Fig. 1B is achieved by using port 1 of the routers 14A, 14B, and auxiliary TNet links LA, to connect to the routers 14A', 14B' of sub-processor systems 10A', 10B'. The topology thereby obtained establishes redundant communication paths between any CPU 12 (12A, 12B, 12A', 12B') and any I/O packet interface 16 of the processing system 10 shown in Fig. 1B. For example, the CPU 12A' of the sub-processor system 10A' may access the I/O 16A of sub-processor system 10A by a first path formed by the router 14A' (in port 4, out... ..shown in Fig. 1B. By interconnecting one port of each router 14 of each sub-processor pair, and using additional auxiliary TNet links LA (illustrated in Fig. 1C with the dotted line connections) between the ports 1 of the routers 14 (14A'' and 14B'') of sub-processor systems 10A'', 10B'' and 10A''', 10B''', two separate, independent data paths can be found between any CPU 12 and any I/O packet interface 16. In this fashion, any end unit (i.e., a CPU 12 or an I/O packet interface 16) will have at least two paths to any other end unit.

Providing alternate paths of access between any two end units (e.g., between a CPU 12 and any other CPU 12, or between any CPU... ..any two of the remaining fault domains. Here, a fault domain could be a sub-processor system (e.g., 10A). Thus, if the sub-processor system 10A were brought down because of a failure the electrical power being supplied, without... ..TNet link LA between the routers 14A''' and 14B''', the CPU 12B of the sub-processor system 10B would have lost access to the I/O packet interface 16''' (via router... ..with the loss of the router 14A (and router 14C) by loss of the sub-processor system 10A, communications between the CPU 12B is still possible via the route of router... ..equally to CPU 12B. As Fig. 2 shows, the CPU 12A includes a pair of processor units 20a, 20b that are configured for synchronized, lock-step operation in that both processor units 20a, 20b receive and execute identical instructions, and issue identical data and command outputs, at substantially the same moments in time. Each of the processor units 20a and 20b is connected, by a bus 21 (21a, 21b) to a corresponding cache memory 22. The particular type of processor units used could contain sufficient internal cache memory so that the cache memory 22 would not... ..22 could be used to supplement any cache memory that may be internal to the processor units 20. In any event, if the cache memory 22 is used, the bus 21 is... ..22 address bits, 3 bits of parity covering the address, and 7 control bits.

The processors 20a, 20b are also respectively coupled, via a separate 64-bit address/data bus 23 to X and Y interface units 24a, 24b. If desired, the address/data communicated on each bus 23a, 23b could also be protected by parity, although this will increase the width of the bus. (Preferably, the processors 20 are constructed to include RISC

R4000 type microprocessors, such as are available from the MIPS Division of Silicon Graphics, Inc. of Santa Clara, California.)

The X and Y interface units 24a, 24b operate to communicate data and command signals between the processor units 20a, 20b and a memory system of the CPU 12A, comprising a memory controller (MC... ..MC halves 26a and 26b) and a dynamic random access memory array 28. The interface units 24 interconnect to each other and to the Mcs 26a, 26b by a 72-bit... ..accompanied by 8 bits of ECC) are written to the memory 28 by the interface units 24, one interface unit 24 will drive only one word (e.g., the 32 most significant portion) of the doubleword being written while the other interface unit 24 writes the other word of the double word (e.g., the least significant 32-bit portion of the doubleword). In addition, on each write operation the interface units 24a, 24b perform a cross-check operation on the data not written by that interface unit 24 with the data written by the other to check for errors; on read operations... ..accessed corresponds to the address of the location from which the doubleword was stored.

Interface units 24a, 24b of the CPU 12A form the circuitry to respectively service the X and Y (I/O) ports of the CPU 12A. Thus, the X interface unit 24a connects by the bi-directional TNet Link Lx to a port of the router 14A of the processor system 10A (Fig. 1A) while the Y interface unit 24b similarly connects to the router 14B of the processor system 10B by TNet Link Ly. The X interface unit 24a handles all I/O traffic between the router 14A and the CPU 12A of the sub-processor system 10A. Likewise, the Y interface unit 24b is responsible for all I/O traffic between the CPU 12A and the router 14B of companion sub-processor system 10B.

The TNet Link Lx connecting the X interface unit 24a to the router 14A (Fig. 1) comprises, as above indicated, two 10-bit buses... ..sub(x)) carries data incoming from the router 14A. In similar fashion, the Y interface unit 24b is connected to the router 14B (of the sub-processor system 10B) by two 10-bit busses: 30(sub(y)) (for outgoing transmissions) and 32... ..y)) (for incoming transmissions), together forming the TNet Link Ly.

The X and Y interface units 24a, 24b are synchronously operated in lock-step, performing substantially the same operations at substantially the same times. Thus, although only the X interface unit 24a actually transmits data onto the bus 30(sub(x)), the same output data is being produced by the Y interface unit 24b, and used for error-checking. The Y interface unit 24b output data is coupled to the X interface unit 24a by a cross-link 34(sub(y)) where it is received by the X interface unit 24a and compared against the same output data produced by the X interface unit. In this way the outgoing data made available at the X port of the CPU... ..the port of the CPU 12A is checked. The output data from the Y interface unit 24b is coupled to the Y port by a 10-bit bus 30(sub(y)), and also to the X interface unit 24a by the 9-bit cross-link 34(sub(y)) where is checked with that produced by the X interface unit.

As mentioned, the two interface units 24a, 24b operate in synchronous, lock-step with one another, each performing substantially the same... ..X and/or Y ports of the CPU 12A must be received by both interface units 24a, 24b to maintain the two interface units in this lock-step mode. Thus, data received by one interface unit 24a, 24b is passed to the other, as indicated by the dotted lines and 9... ..sub(x)) (communicating incoming data being received at the X port by the X interface unit 24a to the Y interface unit 24b) and 36(sub(y)) (communicating data received at the Y port by the Y interface unit 24b to the X interface unit 24a).

Certain more robust operating systems are structured with a fault-tolerant capability in the... ..example, U.S. Patent No. 4,817,091 teaches a multiprocessor system in which each processor periodically messages each of the processors of the system (including itself), under software control, to thereby provide an indication of continuing operation. Each of the processors, in addition to performing its normal tasks, operates as a backup processor to another of the processors. In the event one of the backup processors fails to receive the messaged indication from a sibling processor, it will take over the operation of that sibling (now thought to be inoperative), in... ..platform for both types of software. Thus, when a robust operating system is available, the processing system 10 can be configured to operate in a "simplex" mode in which each of... ..left, in most instances, to software.

Alternatively, for less robust operating systems and software, the processing system 10 provides a hardware-based fault-tolerance by being configured to operate in a... ..g., CPUs 12A, 12B) are coupled together as shown in Fig. 1A, to operate in synchronized, lock-step fashion, executing the same instructions at the substantially the same moment in time... ..data and command symbols. In order to simplify the design of the CPU 12, the processors 20 are precluded from communicating directly with any outside entity (e.g., another CPU 12... ..O device via the I/O packet interface 16). Rather, as will be seen, the processor will construct a data structure in memory and turn over control to the interface units 24. Each interface unit 24 includes a block transfer engine (BTE; Fig. 5) configured to provide a form of... ..to the destination according to information contained in the message packet.

The design of the processing system 10 permits a memory 28 of a CPU to be read or written by... ..via the routers 14. Accordingly, before continuing with the description of the construction of the processing system 10, it would be of advantage to understand first the configuration of the data...information.

As indicated, the HADC message packet operates to communicate write data between the end units (e.g., CPU 12) of the processing system 10. Other message packets, however, may be differently constructed because of their function and... ..CRC. The HC message packet is used to acknowledge a request to write data.

Interface Unit:

The X and Y interface units 24 (i.e., 24a and 24b - Fig. 2) operate to perform three major functions within the CPU 12: to interface the processors 20 to the memory 28; to provide an I/O service that operates transparently to, but under the control of, the processors; and to validate requests for access to the memory 28 from outside sources.

Regarding first the interface function, the X and Y interface units 24a, 24b operate to respectively communicate processors 20a, 20b to the memory controllers (Mcs 26a, 26b) and memory 28 for writing and... ..fast checking of the data read/written. For example, write operations have the two interface units 24a, 24b cooperating to cross-check the data to be written to ensure its integrity (and at the same time, the interface units 24 will operate) to develop an error correcting code (ECC) that covers, as will be appropriate address.

With respect to I/O access, the processors 20 are not provided with the ability to communicate directly with the input/output systems... ..must write data structures to the memory 28 and then pass control to the interface units 24 which perform a direct memory access (DMA) operation to retrieve those data structures, and... ..indicated in the data structure itself.)

The third function of the X and Y interface units 24, access validation to the memory 28, uses an address validation and translation (AVT) table maintained by the interface units. The AVT table contains an address for each system component (e.g., an I/O... ..the incoming message packets are virtual addresses. These virtual addresses are translated by the interface unit to physical addresses recognizable by the memory control units 26 for accessing the memory 28.

Referring to Fig. 5, illustrated is a simplified block diagram of the X interface unit 24a of the CPU 12A. The companion Y interface unit 24b (as well as the interface units 24 of the CPU 12B, or any other CPU 12) is of substantially identical construction. Accordingly, it will be understood that a description of the interface unit 24a will apply equally to the other interface units 24 of the processing system 10.

As Fig. 5 illustrates, the X interface unit 24a includes a processor interface 60, a memory interface 70, interrupt logic 86, a block transfer engine (BTE) 88, access validation and translation logic 90, a packet transmitter 94, and a packet receiver 96.

Processor Interface:

The processor interface 60 handles the information flow (data and commands) between the processor 20a and the X interface unit 24a. A processor bus 23, including a 64 bit address and data bus (SysAD) 23a and a 9 bit command bus 23b, couples the processor 20a and the processor interface 60 to one another. While the SysAD bus 23a carries

memory address and data... ..and qualifying commands carried at substantially the same time on the SysAD bus 23a. The processor interface 60 operates to interpret commands issued by the processor unit 20a in order to pass reads/writes to memory or control registers of the processor interface. In addition, the processor interface 60 contains temporary storage (not shown) for buffering addresses and data for access to... ..26). Data and command information read from memory is similarly buffered en route to the processor unit 20a, and made available when the processor unit is ready to accept it. Further, the processor interface 60 will operate to generate the necessary interrupt signalling for the X interface unit 24a.

The processor interface 60 is connected to a memory interface 70 and to configuration registers 74 by a bi-directional 64 bit processor address/data bus 76. The configuration registers 74 are a symbolic representation of the various control registers contained in other components of the X interface unit 24a, and will be discussed when those particular components are discussed. However, although not specifically... ..throughout other of the logic that is used to implement the X interface 24a, the

processor address/data bus 76 is likewise coupled to read or write to those registers.

Configuration registers 74 are read/write accessible to the processor 20a; they allow the X interface unit to be "personalized." For example, one register identifies the node address of the CPU 12A... ..with the CPU 12A; another, readable only, contains a fixed identification number of the interface unit 24, and still other registers define areas of memory that can be used by, for... ..logic 90, etc.) employing them are discussed.

The memory interface 70 couples the X interface unit 24a to the memory controllers 26 (and to the Y interface unit 24b; see fig. 2) by a bus 25 that includes two 36 bi-directional bit... ..25a, 25b. The memory interface operates to arbitrate between requests for memory access from the processor unit 20, the BTE 88, and the AVT logic 90. In addition to memory accesses from the processor unit 20a, the memory 28 may also be accessed by components of the processing system 10 to, for example, store data requested to be read by the processor unit 20a from an I/O unit 17, or memory 28 may also be accessed for I/O data structures previously set up in memory by the processor unit. Since these accesses are all asynchronous, they must be arbitrated, and the memory interface 70... ..command information accessed from the memory 28 is coupled from the memory interface to the processor interface 60 by a memory read bus 82, as well as to an interrupt logic... ..doubleword quantities. However, while the memory interfaces 70 of both the X and Y interface units 24a ...by the memory interface 70 are coupled to the memory interface by the companion interface unit 24 where they are compared with the same 32 bits for error.

Digressing for the... ..containing interrupt information are received, that information is conveyed to the interrupt logic 86 for processing and posting for action by the processor 20, along with any interrupts generated internal to the CPU 12A. Internally generated interrupts will... ..register 71 (internal to the interrupt logic 86), indicating the cause of the interrupt. The processor 20 can then read and act upon the interrupt. The interrupt logic is discussed more fully below.

The BTE 88 of the X interface unit 24a operates to perform direct memory accesses, and provides the mechanism that allows the processors 20 to access external resources. The BTE 88 can be set-up by the processors 20 to generate I/O requests, transparent to the processors 20 and notify the processors when the requests are complete. The BTE logic 88 is discussed further below.

Requests for... ..8 byte wide format necessary for storing in the memory 28.

Outgoing message packets containing processor originated transaction requests (e.g., a read request asking for a block data from an I/O unit) are monitored by the request transaction logic (RTL) 100. The RTL 100 provides a time... ..will generate an interrupt (handled and reported by the interrupt logic 86) to inform the processor 20 that the request was not honored. In addition, the RTL 100 will validate responses... ..28 (by the DMA operation of the BTE 86) at a location known to the processor 20 so that it can locate the response.

Each of the CPUs 12 are checked... ..discussed. One such check is an on-going monitor of the operation of the interface units 24a, 24b of each CPU. Since the interface units 24a, 24b operate in lock-step synchronism checking can be performed by monitoring the operating states of the paired interface units 24a, 24b by a continuous comparison of certain of their internal states. This approach is implemented by using one stage of a state machine (not shown) contained in the unit 24a of CPU 12A, and comparing each state assumed by that stage with its identical state machine stage in the interface unit 24b. All units of the interface units 24 use state machines to control their operations. Preferably, therefore, a state machine of the memory interface 70 that controls the data transfers between the interface unit 24 and the MC 26 is used. Thus, a selected stage of the state machine used in the memory interface 70 of the interface unit 24a is selected. An identical stage of a state machine of one of the interface unit 24b is also selected. The two selected stages are communicated between the interface units 24a, 24b and received by a compare circuit contained in both interface units 24a, 24b. As the interface units operate lock-step with one another, the state machines will likewise march through the same identical states, assuming each state at substantially the same moments in time. If an interface unit encounters an error, or fails, that activity will cause the interface units to diverge, and the state machines will assume different states. The time will come when... ..that will bring to the attention of the CPUs 12A (or 12B) that the interface units 24a, 24b of that CPU are no longer in lock-step, and to act accordingly... ..X port, receiving only those message packets transmitted by the router 14A of the sub-processor system 10A (Fig. 1A). The Y port is serviced by the Y interface unit 24b to receive message packets from the router 14B of the companion sub-processor system 10B. However, both interfaces (as well as Mcs 26 and processor 20), as has been indicated, are basically mirror images of one another in that both... ..in both structure and function. For this reason, message packet information, received by one interface unit (e.g., 24a) must be passed for processing also to the companion interface unit (e.g., 24b). Further, since both interface units 24a, 24b will assemble the same message packets for transmission from the X or the Y ports, the message packet being transmitted by the interface unit (e.g., 24b) actually being communicated from the associated port (e.g., the Y port) will also be coupled to the other interface unit (e.g., 24a) for cross-checking for errors. These features are illustrated in Figs. 6... ..receiving portions of the packet receivers 96 (96x, 96y) of the X and Y interface units 24a, 24b are broadly illustrated. As shown, each packet receiver 96x, 96y has a clock... ..receive a corresponding one of the TNet Links 32. The CS FIFOs 102 operate to synchronize the incoming command/data symbols to the local clock of the packet receiver 96, buffering... ..104x, coupled to the MUX 104y of the packet receiver 96y of the Y interface unit 24b by the cross-link connection 36(sub(x)). In similar fashion, information received at the Y port is coupled to the X interface unit 24a by the cross-link connection 36(sub(y)). In this manner, the command/data... ..packets received at one of the X, Y ports by the corresponding X, Y, interface unit 24a, 24b is passed to the other so that both will process and communicate the same information on to other components of the interface units 24 and/or memory 28.

Continuing with Fig. 6, depending upon which port X, Y... ..or the other of the CS FIFOs 102x, 102y for communication to the storage and processing logic 110 of the interface unit 24. The information contained in each 9-bit symbol is an 8-bit byte of... ..the encoding of which is discussed below with respect to Fig. 9. The storage and processing logic 110 will first translate the 9-bit symbols to 8-bit data or command... ..the outputs of the CS FIFOs 102x, 102y are also coupled to a command decode unit in addition to the MUX 104. The command decode unit operates to recognize command symbols (differentiating them from data symbols in a manner that is... ..below), decoding them to generate therefrom command signals that are applied to a receiver control unit, a state machine-based element that functions to control packet receiver operations.

As indicated above... ..at the output of the MUX 104, the receiver control portion of the storage control unit enables CRC check logic 106 to calculate a CRC symbol while the data symbols are... ..below, CS FIFOs are found not only in the packet receivers 96 of the interface units 24, but also at each receiving port of the routers 14 and the I/O... ..an even more important part, and perform a unique function, when a pair of sub-processor systems are operating in duplex mode and the two CPUs 12A and 12B of the sub-processor systems 10A, 10B operate in synchronized, lock-step, executing the same instructions at the same time. When operating in this latter... ..difficult to ensure that the clocking regime of the routers 14A and 14B are exactly synchronized to those of the CPUs 12A and 12B - even

when using frequency locked clocking. In... ..used to transmit symbols to a CPU 12 and the clock used by an interface unit 24 to receive those symbols.

The structure of the CS FIFO 102 is diagrammatically illustrated... ..i.e., a packet) or IDLE symbols - except during certain situations (e.g., reset, initialization, synchronization and others discussed below). As explained above, each symbol held in the transmit register 120...same symbol leaving the storage queue, allowing each symbol entering the storage queue 126 to settle before it is clocked out and passed to the storage and processing units 110x (and 110y) by the MUX 104x (and 104y). Since the transmit and receive clocks... ..functioning in duplex mode) operate to transmit symbols with near frequency clocking. Even so, clock synchronization FIFOs are used at these other ports to receive symbols transmitted with near frequency clocking, and the structure of these clock synchronization FIFOs are substantially the same as that used in frequency locked environments, i.e., that... ..of the storage queue 126 are nine bits wide; in near frequency environments, the clock synchronization FIFOs use symbol locations of the queue 126 that are 10 bits wide, the extra... ..the faster clock source. To handle this clock drift, the two pointers are effectively re-synchronized periodically.

When the CPUs 12 are paired and operating in duplex mode, all four interface

units 24 operate in lock-step to, among other things, transmit the same data and receive... ..simplex mode, each independent of the other, clocking need only be near frequency.

The interface unit 24 receives a SYNC CLK signal that is used in combination with a SYNC command symbol to initialize and synchronize the Rcv register 124 to the transmitting router 14. When using either near frequency or... ..102X preferably begin from some known state. Incoming symbols are examined by the storage and processing units 110 of the packet receivers 96. The storage and processing units look for, and act upon as appropriate, command symbols. Pertinent here is that when the... ..receives a SYNC command symbol it will be decoded and detected by the storage and processing unit 110. Detection of the SYNC command symbol by the storage and processing unit 110 causes assertion of a RESET signal. The RESET signal, under synchronous control of the SYNC CLK signal, is used to reset the input buffers (including the clock synchronization buffers) to predetermined states, and synchronize them to the routers 14.

The synchronization of the CS FIFOs 102 of the interface units 24 those ...one or both routers 14A, 14B is discussed more fully below in the section discussing synchronization.

Packet Transmitter:

Each interface unit 24 is assigned to transmit from and receive at only one of the X or Y ports of the CPU 12. When one of the interface units 24 transmits, the other operates to check the data being transmitted. This is an important... ..shows, in abbreviated form, the packet transmitters 94x, 94y of the X and Y interface units 24a, 24b, respectively. Both packet transmitters are identically constructed, so that discussion of one (packet... ..logic 152 that receives, from the BTE 88 or AVT 90 of the associated interface unit (here, the X interface unit 24a) the data to be transmitted - in doubleword (64-bit) format. The packet assembly logic... ..and Y ports: they are either symbols that make up a message packet in the process of being transmitted, or IDLE symbols, or other command symbols used to perform control functions... ..154, 156. The output of the multiplexer 154 connects to the X port. (The interface unit 24b connects the output of the multiplexer 154 to the Y port.) The multiplexer 156... ..sub(x)) to the checker logic 160 of the packet transmitter 94y (of the interface unit 24b).

A selection (S) input of the multiplexers receives a 1-bit output from an... ..is accessible to the MP 18 via an OLAP (not shown) formed in the interface unit 24, and is written with information that "personalizes," among other things, the interface units 24 Here, the X/Y stage of the configuration register 162 configures the packet transmitter 94x of the X interface unit 24a to communicate the X encoder 150x output to the X port; the output of... ..traffic is present, the operation of the two packet interfaces 94 (and, thereby, the interface units 24 with which they are associated) are

continually monitored. Should one of the checkers detect... will be asserted, resulting in an internal interrupt being posted for appropriate action by the processors 20.

Message packet traffic operates in the same manner. Assume, for the moment, that the... that information, a byte at a time, to the X encoder 150x of both interface units 96, which will translate each byte to encoded 9-bit form. The output of the... is checked with that from the packet transmitter 94x. Again, the operation of the interface units 24a, 24b, and the packet transmitters they contain, are inspected for error.

In the same... monitored.

Returning for the moment to Fig. 5, if the outgoing message packet is a processor initiated transaction (e.g., a read request), the processors 20 will expect a message packet to be returned in response. Thus, when the BTE... will issue a timeout signal to the interrupt logic (Fig. 14A) to thereby notify the processors 20 of the absence of a response to a particular transaction (e.g., a read... the access, to name just a few. Also, the area of memory of the memory unit 28 desired to be accessed are identified in the message packets by virtual or I... virtual addresses be translated to physical addresses of the memory 28. Finally, interrupts generated by units or elements external to the CPU 12A, are transmitted via message packets to interrupt the processors 20, which are also written to memory 28 when received. All ...this is handled by the interrupt logic and AVT logic 86, 90.

The AVT logic unit 90 utilizes a table (maintained by the processor 20 in memory 28) containing AVT entries for each possible external source permitted access to the memory 28. Each AVT entry identifies a specific source element or unit and the particular page (a page being nominally 4K (4096) bytes), or portion of a... expected" memory accesses. Expected memory accesses are those initiated by the CPU 12 (i.e., processors 20) such as a read request for information from an I/O device. These latter memory accesses are handled by a transaction sequence number (TSN) assigned to each processor initiated request. At about the time the read request is generated, the processors 20 will allocate an area of memory for the data expected to be received in... and 26b are, in turn, respectively coupled to the memory interfaces 70 of each interface unit 24a, 24b. The 64-bit doublewords are written to the memory 28 with the upper... check bits respectively from the memory interfaces 70 (70a, 70b) of each of the interface units 24a, 24b (Fig. 5).

Referring to Fig. 10, each memory interface 70 receives, from either the bus 82 from the processor interface 60 or the bus 83 from AVT logic 90 (see Fig. 5), of the associated interface unit 24, 64 bits of data to be written to memory. The busses 76 and 83... other for cross-checking between them. Thus, for example, the memory interface 70a (of interface unit 24a) will drive the MC 26a with the "upper" 32 bits of the 64 bits... are check bits, leaving 40 bits unused.

Access Validation:

As previously indicated, components of the processing system 10 external to the CPU 12A (e.g., devices of the I/O packet... not without qualification. Access validation, as implemented by the AVT logic 90 of the interface units 24, operates to prevent the content of the memory 28 from being ...Accesses to the memory 28 are validated by the AVT logic 90 of each interface unit 24 (Fig. 5), using all of six checks: (1) that the CRC of the message... also are permitted the particular message packet source.

The access validation mechanism of the interface unit 24a, AVT logic 88, is shown in greater detail in Fig. 11. Incoming message packets...and post an interrupt to the interrupt logic 86 (Fig. 5) for action by the processor 20.

The mask operation permits the size of the table of AVT entries to be varied. The content of the AVT mask register 175 is accessible to the processor 20, permitting the processors 20 to optionally select the size of the AVT entry table. A maximum AVT table... 172 allows the AVT size to be matched to the needs of the system. A processing system 10 that includes a larger number of external elements (e.g., the number of... amount of the memory space of memory 28 to the AVT entries. Conversely, a smaller processing system 10, with a smaller number of external elements will not have such a large... set to a logic "ZERO" indicate an nonexistent TNet address, outside the

limits of the processing system 10. A received packet with a TNet address outside the allowable TNet range will...
...in Fig. 11 as being held in the AVT entry register 180 during the validation process. AVT entries have two basic formats: normal and interrupt. The format of a normal AVT...of the AVT input register 170) will result in an error being posted to the processor via an interrupt.

A 12-bit "Permissions" field is included in the AVT entry to...path=0). Denials are logged as interrupts with the interrupt logic, and reported to the processor 20 - if the E field is set to a state ("ONE") that enables error-reporting...e.g., to a "ONE"), the other fields (Upper Bound, etc.) gain new definitions for processing interrupt writes and managing interrupt queues. This is discussed in more detail below in connection...memory 28 will be handled. Set to one state, the requested write operation will be processed normally; set to a second state, write requests specifying addresses with a fractional cache line...be written to a specific queue (interrupt queue) in memory 28, with signalling provided the processors 20 to indicate that an interrupt has been received and "posted," and ready for servicing by the processors 20. Since the interrupt queues are at specific memory locations, the processor can obtain the interrupt data when needed.

An AVT interrupt entry for an interrupt may...by the interrupt logic 86, and extracted from the head of the queue by the processor 20 when servicing the interrupt.

The AVT interrupt entry also includes a 20-bit segment ("Source ID") containing source ID information, identifying the external unit seeking attention by the interrupt process. If the source ID information of the AVT interrupt entry does not match that contained...class" of the interrupt that is used to determine the interrupt level set in the processor 20 (described more fully below); (2) a queue number that is used to select, as...capability to deliver interrupts to a CPU 12 for servicing. For example, an I/O unit may be unable to complete a read or write transaction issued by a CPU because...identify the recipient. These and other errors, exceptions, and irregularities, noted by the I/O units, or the I/O Interface elements, can become the a condition that requires the intervention...the AVT entry register 180 for use by the interrupt logic 86 of the interface unit 24 (Fig. 5), illustrated in greater detail in Fig. 14A.

It is interrupt logic 86...four circular queues specified by the base address information contained in the AVT entry. The processor(s) 20 will then be notified, and it will be up to them as to...selected tail queue register 256 by combiner circuit 270, the output of which is the processed by the "mod z" circuit 273 to turn new offset into the queue at which...signal. The Queue Full warning signal becomes an "intrinsic" interrupt that is conveyed to the processor units 20 as a warning that if the matter is not promptly handled, later-received interrupt...will be discarded.

Incoming message packet interrupts will cause interrupts to be posted to the processor 20 by first setting one of a number of bit positions of an interrupt register 280. Multi-entry queued interrupts are set in interrupt registers 280a for posting to the processor 20; single-entry queue interrupts use interrupt register 280b. Which bit is set depends upon...multi-entry queued interrupts, soon after a multi-entry queued interrupt is determined, the interface unit will assert a corresponding interrupt signal (I1) that is applied to decode circuit 283. Decode...of register 280a to set, thereby providing advance information concerning the received interrupt to the processor(s) 20, i.e., (1) the type of interrupt posted, and (2) the class of...to one another by a compare circuit 279. The update register is writable by the processor 20 to select a register pair for comparison. If the content of the two selected...cleared.

Digressing for the moment, there are two basic types of interrupts that concern the processors 20: those interrupts that are communicated to the CPU 12 by message packets, and those...the seven interrupt postings to a latch 288, from which they are coupled to the processor 20 (20a,20b) which has an interrupt register for receiving holding the postings.

In addition... ..change in interrupts (either an interrupt has been serviced, and its posting deleted by the processor 20, or a new interrupt has been posted), a "CHANGE" signal will be issued to the processor interface 60 to inform it that an interrupt posting change has occurred, and that it should communicate the change to the processor 20.

Preferably, the AVT entry register 180 is configured to operate like a single line... ..such as set-associative, fully-associate, or direct-mapped, to name a few.

Coherency:

Data processing systems that use cache memory have long recognized the problem of coherency: making sure that... ..the incoming packet is permitted access are applied to a boundary crossing (Bdry Xing) check unit 219. Boundary check unit 219 also receives an indication of the size of the cache block the CPU 12... ..Len field of the header information from the AVT input register 170. The Bdry Xing unit determines if the data of the incoming packet is not aligned on a cache boundary...time an interrupt will be written to the queued interrupt register 280, to alert the processors 20 that a portion of the incoming data is located in the special queue.

In... ..not, the packet (both header and data) is written to a special queue, and the processors so notified by the intrinsic interrupt process described above. The processors may then move the data from the special queue to cache 22, and later write... ..the cache 22 and the memory 28 is preserved.

Block Transfer Engine (BTE):

Since the processor 20 is inhibited from directly communicating (i.e., sending) information to elements external to the... ..indirect method of information transmission.

The BTE 88 is the mechanism used to implement all processor initiated I/O traffic to transfer blocks of information. The BTE 88 allows creation of... ..BTE registers 300, 302 whose content is coupled to the MUX 306 (of the interface unit 24a; Fig. 5) and used to access the system memory 28 via the memory controllers... ..BTE data structure 304 in the memory 28 of the CPU 12A (Fig. 2). The processors 20 will write a data structure 304 to the memory 28 each time information is... ..begin on a quadword boundary, and the BTE registers 300, 302 are writable by the processors 20 only. When a processor does write one of the BTE registers 300, 302, it does so with a word... ..the request bit (rc0, rcl) to a clear state, which operates to initiate the BTE process, which is controlled by the BTE state machine 307.

The BTE registers 300, 302 also... ..cause (ec) bit differentiates time-outs and NAKs.

When information is being transferred by the processors 20 to an external unit, the data buffer portion 304b of the data structure 304 holds the information to be transferred. When information from an external unit is received by the processors 20, the data buffer portion 304b is the location targeted to hold the read response information.

The beginning of the data structure 304, portion 304a written by the processor 20, includes an information field (Dest), identifying the external element which will receive the packet... ..the transmitted data is to be written. This information is used by the packet transmitter unit 120 (Fig. 5) to assemble the packet in the form shown in Figs. 3-4...list (el) bit, when set, indicates the end of the chain, and halts the BTE processing.

The interrupt completion (ic) bit, when set, will cause the interface unit 24a to assert an interrupt (BTECmp) which sets a bit in the interrupt register 280... ..the chain pointer).

The interrupt time-out (it) bit, when set, will cause the interface unit 24a to assert an interrupt signal for the processor 20 if the acknowledgement of the access times-out (i.e., if the request timer... ..time), or elicits a NAK response (indicating that the target of the request could not process the request).

Finally, if the check sum (cs) bit is set, the data to be... ..containing the data from which the check sum was formed.

To sum up, when the processors 20 of the CPU 12A desire to send data to an external unit, they will write a data structure 304 to the memory 28, comprising identifier information in portion 304a of the data structure, and the data in the buffer portion 304b. The processors 20 will then determine the priority of the data and will write the BTE register... information, and sent.

If the data structure 304 indicates a read request (i.e., the processors 20 are seeking data from an external unit - either an I/O device or a CPU 12), the Len and Local Buffer Ptr... receiver 100 (Fig. 5) until the local memory write operation is executed.

Responses to a processor-generated read request to an external unit are not processed by the AVT table logic 146. Rather, when the processors 20 set up the BTE data structure, a transaction sequence number (TSN) is assigned the... the BTE 88, which will be an HAC type packet (Fig. 4) discussed above. The processors 20 will also include an memory address in the BTE data structure at which the...302, assume that the foregoing transfer of data from the CPU 12A to an external unit is of a large block of information. Accordingly, a number of data structures would be set up in memory 28 by the processors 20, each (except the last) including a chain pointer to additional data structures, the sum... sent. Assume now that a higher priority request is desired to be made by the processors 20. In such a case, the associated data structure 304 for such higher priority request... with another BTE operation descriptor.

Memory Controller:

Returning, for the moment, to Fig. 2, interface units 24a, 24b access the memory 28 via a pair of memory controllers (MC) 26a, 26b. The Mcs provide a fail-fast interface between the interface units 24 and the memory 28. The Mcs provide the control logic necessary for accessing... in dynamic random access memory (DRAM) logic). The Mcs receive memory requests from the interface units 24, and execute reads and writes as well as providing refresh signals to the DRAMs... to provide a 72 bit data path between the memory array 28 and the interface units 24a, 24b, which utilize an SBC-DBD-SbD ECC scheme, where $b=4$, on a... 26a, 26b to work together and simultaneously supply a 64-bit word to the interface units 24 with minimum latency, one-half of which (DO) comes from the MC 26a, and the other half (D1) comes from the other MC 26b. The interface unit 24 generate and check the ECC check bits. The ECC scheme used will not only... 26 bus 25, as well as in internal registers.

From the viewpoint of the interface units 24, the memory 28 is accessed with two instructions: a "read N doubleword" and a... doubleword read or a block read format. The signal called "data valid" tells the interface units 24 two cycles ahead of time that read data is being returned or not being returned.

As indicated above, the maintenance processor (MP 18; Fig. 1A) has two means of access to the CPUs 12. One is... 18 will write a register contained in the OLAP 285 with instructions that permit the processors 20 to build an image of a sequence of instructions in the memory that will permit them (the processors 20) to ...to transfer instructions and data from an external (storage) device that will complete the boot process.

The OLAP 285 is also used by the processors 20 to communicate to the MP 18 error indications. For example, if one of the interface units 24 detect a parity error in data received from the memory controller 26, it will... and address transfers on the bus 25 between the MC 26a and the corresponding interface unit 24a. The addressing and data transfers on the DRAM data bus, as well as generation... the CPU 12.

Packet Routing:

The message packets communicated between the various elements of the processing system 10 (e.g., CPUs 12A, 12B, and devices coupled to the I/O packet... First, each TNet Link L connects to an element (e.g., router 14A) of the processing system 10 via a port that has both receive and transmit capability. Each transmit port... cycle (i.e., each clock period) of the T(underscore)Clk so that the clock

synchronization FIFO at the receiving end of the transmission will maintain synchronization.

Clock synchronization is dependent upon the mode in which the processing system 10 is operated. If operating in the simplex mode in which the CPUs 12A... connect directly to the CPUs may drift with respect to each other. Conversely, when the processing system 10 operates in a duplex mode (e.g., the CPUs operate in synchronized, lock-step operation), the clocks between routers 14 and the CPUs 12 to which they... not necessarily phase-locked).

The flow of data packets between the various elements of the processing system 10 is controlled by command symbols, which may appear at any time, even within... initiated by a CPU 12, or MP 18, and promulgated to all elements of the processing system 10 by the routers 14 to communicate an event requiring software action by all... command symbol is used in conjunction with near frequency operation as an aid to maintaining synchronization between the two clock signals that (1) transfer each symbol to, and load it in each receiving clock synchronization FIFO, and (2) that retrieves symbols from the FIFO.

SLEEP: This command symbol is sent by any element of the processing system 10 to indicate that no additional packet (after the one currently being transmitted, if... received.

SOFT RESET (SRST): The SRST command symbol is used as a trigger during the processes ("synchronization" and "reintegration," described below) that are used to synchronize symbol transfers between the CPUs 12 and the routers 14A, 14B, and then to place... SYNC command symbol is sent by a router 14 to the CPU 12 of the processing system 10 (i.e., the sub-processor systems 10A/10B) to establish frequency-lock synchronization between CPUs 12 and routers 14A, 14B prior to entering duplex mode, or when in duplex mode to request synchronization, as will be discussed more fully below. The SYNC command symbol is used in conjunction... or duplex to simplex), among other things, as discussed further below in the section on Synchronization and Reintegration.

THIS LINK BAD (TLB): When any system element receiving a symbol from a TNet link L (e.g., a router, a CPU, or an I/O unit) notes an error when receiving a command symbol or packet, it will send a TLB... identical pairs of symbols that are compared to one another when pulled from the clock synchronization FIFOs. The DVRG command symbol signals the CPU 12 that a mis-compare has been noted. When received by the CPUs, a divergence detection process is entered whereby a determination is made by the CPUs which CPU may be failing... command symbols described above operate to control message flow between the various elements of the processing system 10 (e.g., CPUs 12, router 14, and the like), using principally the BUSY however, an "end node" (i.e., a CPU 12 or I/O unit 17 - Fig. 1) may not assert backpressure because one of its transmit ports is backpressured... Improperly addressed packets are discarded by the router 14.

When a system element of the processing system 10 receives a BUSY command symbol on a TNet link L on which it... other command symbols (READY, BUSY, etc.).

Whenever a TNet port of an element of the processing system 10 detects receipt of a READY command symbol, it will terminate transmission of FILL... receives.

As will be seen, all elements (e.g., router 14, CPUs 12) of the processing system 10 that connect to a TNet link L for receiving transmitted symbols will receive those symbols via a clock synchronization (CS) FIFO. For example, as discussed above, the interface units 24 of CPUs 12 include all CS FIFOs 102x, 102y (illustrated in Fig. 6). The... depth to allow for speed matching, and the elastic FIFOs must provide sufficient depth for processing delays that may occur between transmission of a BUSY command symbol during receipt of a... another data byte in packet B. As packet A progresses to the next router, the process would be repeated. If the router 14 displaces more data bytes than the FIFO can... irrespective of its own findings.

SLEEP Protocol:

The SLEEP protocol is initiated by a maintenance processor via a maintenance interface (an on-line access port - OLAP), described below. The SLEEP protocol... ..reintegrate a slice of the system 10. Routers 14 must be idle (no packets in process) in order to change modes without causing data loss or corruption. When a SLEEP command symbol is received, the receiving element of processing system 10 inhibits initiation of transmission of any new packet on the associated transmit port... ..The HALT command symbol provides a mechanism for quickly informing all CPUs 12 in a processing system 10 that is necessary to terminate I/O activity (i.e., message transmissions between... ..CPUs that receive HALT command symbols on either of their receive ports (of the interface units 24) will post an interrupt to the interrupt register 280 if the system halt interrupt... ..interrupt; Fig. 14A).

The CPUs 12 may be provided with the ability to disable HALT processing. Thus, for example, the configuration registers 75 of the interface units 24 can include a "halt enable register" that, when set to a predetermined state (eg., ZERO) disables HALT processing, but reporting detection of a HALT symbol as an error.

Router Architecture:

Referring now to... ..simplified block diagram of the router 14A is illustrated. The other routers 14 of the processing system 10 (e.g., routers 14B, 14', etc.) are of substantially identical construction and, therefore...these ports 4, 5 are structured to operate in a frequency locked environment when a processing system 10 is set for duplex mode operation. In addition, when in duplex mode, a... ..5)) will receive the command/data symbols from the CPUs, pass them through the clock synchronization FIFOs 518 (discussed further below), and compare each symbol exiting the clock synchronization FIFOs with a gated compare circuit 517. When duplex operation is entered, a configuration register... ..517 to activate the symbol by symbol comparison of the symbols emanating from the two synchronization FIFOs 518 of the router input logic 502 for the ports 4 and 5. Of... ..to that received, at substantially the same time, by the other port input.

To maintain synchronization in the duplex mode, the two port outputs of the router 14A that transmit to... ..mode, are duplicated by the routers 14, and returned to both CPUs.) The output logic units 504(sub(4)), 504(sub(5)) that are coupled directly to the CPUs 12 will... ..message packet identifies only one of the duplexed CPUs 12, e.g., CPU 12A) in synchronized fashion, presenting those symbols in substantially simultaneous fashion to the two CPUs 12. Of course, the CPUs 12 (more accurately, the associated interface units 24) receive the transmitted symbols with synchronizing FIFOs of substantially the same structure as that illustrated in Fig. 7A so that, even... ..from the FIFO structures by both CPUs 12 on the same instruction cycle, maintaining the synchronized, lock-step operation of the CPUs 12 required by the duplex operating mode.

As will... ..conjunction with configuration data written to registers contained in control logic 509 by the maintenance processor 18 (via the on-line access port 285' and serial bus 19A; see Fig. 1A...links L. The input logic 505 of each port input 502 also assists in maintaining synchronization - at least for those ports sending symbols in the near-frequency environment - by removing received... ..slower-receiving element receiving symbols from a faster-sending element could overload the input clock synchronization FIFO of the slower-receiving element. That is, if a slower clock is used to pull symbols from the clock synchronization FIFO put there by a faster clock, ultimately the clock synchronization FIFO will overflow.

The preferred technique employed here is to periodically insert SKIP symbols in... ..stream to avoid, or at least minimize, the possibility of an overflow of the clock synchronization FIFO (i.e., clock synchronization FIFO 518; Fig. 20A) of a router 14 (or CPU 12) due to a T... ..being slightly higher in frequency than the local clock used to pull symbols from the synchronization FIFO. Using SKIP symbols to by-pass a push (onto the FIFO) operation has the... ..stall each time a SKIP command symbol is received so that, insofar as the clock synchronization FIFO is concerned, the transmitting clock that accompanied the SKIP symbol was missing.

Thus, logic... ..the port inputs 502 will recognize, and key off receipt of, SKIP command symbols for synchronization in the near frequency clocking environment so that nothing is pushed onto the FIFO, but... ..14, or between routers 14, or between a router 14 and an I/O interface unit 16A - Fig. 1) at a 50 Mhz rate, this allows for a

worst case frequency... ..symbol by supplying FILL or IDLE symbols (which are received and pushed onto the clock synchronization FIFOs, but are not passed to the elastic FIFOs). In short, each elastic FIFO 506...received symbols are then communicated from the input register 516 and applied to a clock synchronization FIFO 518, also by the T(underscore)Clk. The clock synchronization FIFO 518 is logically the same as that illustrated in Figs. 8A and 8B, used in the interface units 24 of the CPUs 12. Here, as Fig. 20A shows, the clock synchronization FIFO 518 comprises a plurality of registers 520 that receive, in parallel, the output of... ..516. Associated with each of the registers 520 is a two-stage validity (V) bit synchronizer 522, shown in greater detail in Fig. 20B, and discussed below. The content of each... ..registers 520, together with the one-bit content of each associated two-stage validity bit synchronizer 522, are applied to a multiplexer 524, and the selected register/synchronizer pulled from the FIFO, and coupled to the elastic FIFO 506 by a pair of... ..is determined the state of the Push Select signal provided by a push pointer logic

unit 530; and, selection of which register 520 will supply its content, via the MUX 524... ..and loading of the register 520 selected by the push pointer logic 530. Similarly, the synchronization FIFO control logic 534 receives the clock signal local to the router (Rcv Clk) to... ..pointer logic 532.

Digressing for a moment, and referring to Fig. 20B, the validity bit synchronizer 522 is shown in greater detail as including a D-type flip-flop 541 with... ..530 (Fig. 20A) selects the register 520 of the FIFO with which the validity bit synchronizer is associated for receipt of the next symbol - if not a SKIP symbol.

The delay... ..Truth Table, below). The D-type flip-flop 543 acts as an additional stage of synchronization, ensuring a stable level at the V output relative to the local Rec Clk. The... ..flip-flop 542, allowing the Pull signal (a periodic pulse from the sync FIFO Control unit 534) to clear the validity bit on this validity synchronizer 522 when the associated register 520 has been read.

(Table omitted)

In summary, the validity synchronizer 522 operates to assert a "valid" (V) signal when a symbol is loaded in a...blocked from being routed out a particular port because another message is already in the process of being routed out that port. However, that other message in turn is also blocked...an incoming message packet bound for the CPUs will be replicated by the crossbar logic unit by routing the message packet to both port output 504(sub(4)) and 504(sub... ..P) identifies which of path (X or Y) should be used for accessing two sub-processing the device.

The routers 14 provide a capability of constructing a large, versatile routing network for, for example, massively parallel processing architectures. Routers are configured according to their location (i.e., level) in the network by...j)) and 509(sub(k)) are such that bits "def" are used in the algorithmic process, then bits "abc" of the Region ID are compared to the content of the Device... ..the route to default register 509(sub(f))) to the final stage of the selection process: check logic 602. Check logic 602 operates to check the status of the port output...a lower level router, and may be located in one or another of the sub-processing systems 10A, 10B. Whether a router is an upper level or lower level router depends... ..of CPUs 12 and I/O devices 16 to one another, forming a massively parallel processing (MPP) system. Other such MPP systems may exist, and it is those routers configured as... ..captured. As soon as the message packet's Destination ID is so captured, the selection process begins, proceeding to the development of a target port address that will be used to...an error that will be posted to the MP18 via the router's (or interface unit's) OLAP for action.

Digressing, it should be appreciated that these protocol rules observed by the routers 14 are also observed by the CPUs 12 (i.e., interface units 24) and I/O packet interfaces 17.

Finally, when the router 14A is in the... ..directly with the CPUs 12A, 12B, and duplex mode is used, a duplex operation logic unit 638 is utilized to coordinate the port output connected to one of the CPUs 12A... ..was able to write instructions to the OLAP 285 that would be executed by the processors 20 to build a small memory image and

routine to permit the CPU 12 to... ..the clock generation circuit design. There will be one clock generator circuit in each sub-processor system 10A/10B (Fig. 1) to maintain synchronism. Designated generally with the reference numeral 650... ..used by the various elements (e.g. CPU. 12, routers 14, etc.) of the sub-processor system containing the clock circuit 650 (e.g., 10A).

The clock generator 654 is shown...The 50 Mhz clock signals produced by the counter 663 are distributed throughout the sub-processor system where needed.

Turning now to Fig. 25, there is illustrated the interconnection and use... ..the clock circuits 650 used to develop synchronous clock signals for a pair of sub-processor systems 10A, 10B (Fig. 1) for frequency locked operation. As illustrated in Fig. 25, the two CPUs 12A and 12B of the sub-processor systems 10A, 10B each have a clock circuit 650, shown in Fig. 25 as clock... ..654B of both CPUs 12. A driver and signal line 667 interconnects the two sub-processor systems to deliver the M(underscore)CLK signal developed by the oscillator circuit 652A to the clock generator 654B of the sub-processor system 10B. For fault isolation, and to maintain signal quality, the M(underscore)CLK signal is delivered to the clock generator 654A of the sub-processor system 10A through a separate driver and a loopback connection 668. The reason for the... ..way, the cable (not shown) will establish the connection shown in Fig. 25 between the sub-processor systems 10A, 10B; connected another way, the connections will be similar, but the oscillator 652B... ..Fig. 25, the M(underscore)CLK signal produced by the oscillator circuit 652A of sub-processing system 10A is used by both sub-processing systems 10A, 10B as their respective SYNC CLK signals and the various other clock signals... ..produced by the clock generators 654A, 654B. Thereby, the clock signals of the paired sub-processing systems 10A, 10B are synchronized for the frequency locked operation necessary for duplex mode.

The VCXOs 662 of the clock... ..This allows both clock generators 654A, 654B to continue to provide to the two sub-processing systems 10A, 10B clock signals in the face of improper operation of the oscillator circuit 652A, although the sub-processor systems may no longer be frequency-locked.

The LOCK signals asserted by the phase comparators... ..LOCK signal signifies that the 50 Mhz signals produced by a clock generator 654 are synchronized, both in phase and in frequency, to the M(underscore)CLK signal. Thus, if either... ..signal that accompanies the symbol stream, and is used to push symbols onto the clock synchronizing FIFO of the receiving element (router 14, or CPU 12) is substantially identical in frequency... ..not phase, to that of the receiving element used to pull symbols from the clock synchronization FIFOs. For example, referring to Fig. 23, which illustrates symbols being sent from the router... ..clock (Local Clk). The former (Rcv Clk) is used to push symbols onto the clock synchronization FIFOs 126 of each CPU, whereas the latter is used to pull symbols from the... ..much higher frequency clock signal. In such situations provision must be made to ensure that synchronization is maintained between the two CPUs as to symbols pulled from the clock synchronization FIFOs 126 of each.

Here, a constant ratio clocking mechanism is used to control operation of the two clock synchronization FIFOs 126, providing the clock signal that pulls symbols from the two FIFOs at the... ..control mechanism is shown, designated with the reference numeral 70. As Fig. 26A illustrates, clock synchronization FIFO control mechanism 700 includes an pre-settable, multi-stage serial shift register 702, the... ..ratio of the clock signal at which symbols are communicated and pushed onto the clock synchronization FIFOs 126 to the frequency of the clock signal used locally. Here, a 15 stages... ..that will be used as the Local Clk signal to pull symbols from the clock synchronization FIFOs 126, and to operate (update) the pull pointer counter 130. The selected output is... ..of the CPU 12 to the clock signal used to push symbols onto the clock synchronization FIFO 126, Rcv Clk, the serial shift register is preset so that M stages of... ..duplexed CPUs 12 with a 50 Mhz clock. Thus, symbols are pushed onto the clock synchronization FIFOs 126 of the CPUs at a 50 Mhz rate. Assume further that the clock... ..of the MUX 704, which produces the clock signal that pulls symbols from the clock synchronization FIFOs 126, Rcv Clk, will contain, for each 100 ns period, five clock pulses. Thus... ..five symbols will be pushed onto, and five symbols will be pulled from, the clock synchronization FIFOs 126.

This example is symbolically shown in Fig. 26B, while the timing diagram shown... ..labelled "IN" in Fig. 27) of the Rev Clk will push symbols onto the clock synchronization FIFOs 126. During that same 100 ns period, the serial shift register 702 circulates a... ..clocks which would require additional storage (i.e., an increase in the size of the synchronization FIFO) and impose more latency.

The constant ratio clock circuit presented here (Figs. 26) is... ..frequency to a clock regime of a different, higher frequency. The use of a clock synchronization FIFO is necessary here for compensating effects of signal delays when operating in synchronized, duplexed mode to receive pairs of identical command/data symbols from two different sources. However...so long as there are at least two registers in the place of the clock synchronization FIFO. Transferring data from a higher-frequency clock regime to a lower frequency clock regime... ..a wide range of possible clock ratios.

I/O Packet Interface:

Each of the sub-processor systems 10A, 10B, etc. will have some input/output capability, implemented with various peripheral units, although it is conceivable that the I/O of other sub-processor systems would be available so that a sub-processing system may not necessarily have local I/O. In any event, if local I/O... ..device (e.g., a signal line) would be received by the I/O packet interface unit 16 and used to form an interrupt packet that is sent to the CPU 12... ..OLAP bus, configuration information.

On-Line Access Port:

The MP 18 connects to the interface unit 24, memory controller (MC) 26, routers 14, and I/O packet interfaces with interface signals... ..OLAP 258 is essentially the same, regardless of what element (e.g. router 14, interface unit 24, etc.) it is used with. Fig. 28 diagrammatically illustrates the general structure of the... ..circuit chip used to implement certain of the elements discussed herein. For example, each interface

unit 24, memory controller 26, and router 14 is implemented by an application specific integrated circuit... ..of the OLAP 158 shown in Fig. 28 describes the OLAP associated with the interface unit 24, the MC 26, and the router 14 of the system.

As Fig. 28 shows...asymmetric variables, a "soft-vote" (SV) logic element 900 (Fig. 30A) is provided each interface unit 24 of each CPU 12. As Fig. 30 illustrates, the SV logic elements 900 of each interface unit 24 are connected to one another by a 2-bit SV bus 902, comprising bus lines 902a and 902b. Bus lines 902a carry one-bit values from the interface units 24 of CPU 12A to those of CPU 12B. Conversely, bus line 902b carries one... ..the CPU 12A.

Illustrated in Fig. 30B, is the SV logic element 900a of interface unit 24a of CPU 12A. Each SV logic element 900 is substantially identical in construction and... ..900a should be understood as applying equally to the other logic elements 900a (of interface unit 24b, CPU 12A), and 900b (of the interface units 24a, 24b of CPU 12B) unless noted otherwise. As Fig. 30B illustrates, the SV logic interface units 24a, 24b of the CPU 12A can communicate asymmetrical variables to each other.

In a... ..to the remote register 907 of logic element 902a (and that of the other interface unit 24b).

The logic elements 902 form a part of the configuration registers 74 (Fig. 5). Thus, they may be written by the processor unit(s) 20 by communicating the necessary data/address information over at least a portion of... ..local and remote registers 906 and 907.

The MUX 914 operates to provide each interface unit 24 of CPU 12A with selective use of the bus line 902a for the SV logic elements 900a, or for communicating a BUS ERROR signal if encountered during the reintegration process (described below) used to bring a pair of CPUs 12 into lock-step, duplex operation... ..same time, write the enable registers 912 of the logic element 900 of both interface units 24 of each CPU. One of the two logic elements

900 of each CPU will... it is the output enable registers 912 associated with the logic elements 900 of interface units 24a of both CPUs 12A, 12B that are written to enable the associated drivers 916. Thus, the output registers 904 of the interface units 24a of each CPU will be communicated to the bus lines 902; that is, the... to the bus line 902a, while the output register associated with logic element 900b, interface unit 24a of CPU 12B is communicated to bus line 902b. The CPUs 12 will both... again written by each CPU, followed again by reading the remote input registers 907. This process is repeated, one bit at a time, until the entire variable is communicated from the... each CPU 12 to the remote input register of the other. Note that both interface units 24 of CPU 12B will receive the bit of asymmetric information.

One example of use... elements 900 are also used to communicate bus errors that may occur during the reintegration process to be described. When reintegration is being conducted, a REINT signal will be asserted. As... ERROR signal is selected by the MUX 914 and communicated to the bus line 902a.

Synchronization:

Proper operation of the sub-processing systems 10A, 10B (Figs. 1A, 2) whether operating independently (simplex mode), or paired and operating in synchronized lock-step (duplex mode), requires assurance that data communicated between the CPUs 12A, 12B and... routers 14A, 14B will be received properly, and that any initial content of the clock synchronization FIFOs 102 (of CPUs 12A, 12B; Fig. 5) and 519 (of routers 14A, 14B; Fig... erroneously interpreted as data or commands. The push and pull pointers of the various clock synchronization FIFOs 102 (in the CPUs 12) and 518 (in the routers 14) need to be... apart, and presetting the associated FIFO queues to some known state. This done, all clock synchronization FIFOs are initialized for near... in order to properly implement the lock-step operation of duplex mode operation, the clock synchronization FIFOs must be synchronized to operate with the particular source from which they receive data in order accommodate any... 14A, 14B to the CPUs 12A, 12B must be accounted for. It is the clock synchronization FIFOs 102 of the paired CPUs 12 that operate to receive message packet symbols, adjust... and present symbols to the two CPUs in a simultaneous manner to maintain lock-step synchronization necessary for duplex mode operation.

In similar fashion, each symbol received by the routers 14A... the CPUs (which is discussed further hereinafter). Again, it is the function of the clock synchronization FIFOs 518 of the routers 14A, 14B that receive message packets from the CPUs 12... so that the symbols received from the two CPUs 12 are retrieved from the clock synchronization FIFOs simultaneously.

Before discussing how the clock synchronization FIFOs of the CPUs and routers are reset, initialized, and synchronized, an understanding of their operation to maintain synchronous lock-step duplex mode operation is believed helpful. Thus, referring for the moment to Fig. 23, the clock synchronization FIFOs 102 of the CPUs 12A, 12B that receive data, for example, from the router... underscore)Clk, from the router 14A to the CPU 12B.

Consider operation of the clock synchronization FIFOs 102(sub(x)), 102(sub(y)), to receive identical symbol streams during duplex operation... held by the push and pull pointer counters 128, 130 for the CPU 12A (interface unit 24a), and the content of each of the four storage locations (byte 0. ..., byte 3... 6 show the same thing for the FIFO 102(sub(y)) of CPU 12B interface unit 24a for each symbol of the duplicated symbol stream.

Assuming the delay 640 is no...0" locations of the queues 126. This is because (1) the FIFOs 102 have been synchronized to operate in synchronism (a process described below), and (2) the push pointer counters 128 are clocked by the clock signal... of the symbol stream transmitted by the router 14A will be pulled from the clock synchronization FIFOs 102 of the CPUs 12A, 12B simultaneously, maintaining the required synchronization of received data when operating in duplex mode. In effect, the depths of the queues... order to achieve the operation just described with reference to Table 6, the reset and synchronization process shown in Fig 31A is used. The process not only initializes the clock synchronization FIFOs 102 of the CPUs 12A, 12B for duplex mode operation, but also operates to adjust the clock synchronization FIFOs 518 (Fig. 19A) of the CPU ports of each of the routers 14A, 14B for duplex operation. The reset and synchronization process uses the SYNC command symbol

to initiate a time period, delineated by the SYNC CLK signal 970 (Fig. 31B), to reset and initialize the respective clock synchronization FIFOs of the CPUs 12A and 12B and routers 14A, 14B. (The SYNC CLK signal... It is of a lower frequency than that used to receive symbols by the clock synchronization FIFOs, T(underscore)Clk. For example, where T(underscore)Clk is approximately 50 MHz, the... signal is approximately 3.125 MHz.)

Turning now to Fig. 31A, the reset and initialization process begins at step 950 by switching the clock signals used by the CPUs 12A, 12B and routers 14A, 14B as the transmit (T(underscore)Clk) and the unit's local clock (Local Clk) clock signals so that they are derived from the same... In addition, configuration registers in the CPUs 12A, 12B (configuration registers 74 in the interface units 24) and the routers 14A, 14B (contained in control logic unit 509 of routers 14A, 14B) are set to the FreqLock state.

The following discussion involves step 952, and makes reference to the interface unit 24 (Fig. 5), router 14A (Fig. 19A) and Figs. 31A and 31B. With the clock... otherwise be sent followed by a self-addressed message packet. Any message packet in the process of being received and retransmitted when the SLEEP command symbols are received and recognized by... per the destination address). The SLEEP command symbol operates to "quiece" router 14A for the synchronization process. The self-addressed message packet sent by the CPU 12A, when received back by the... message packet sent after the SLEEP command symbol would necessarily have to be the last processed by the router 14A.

At step 954 the CPU 12A checks to see if it...the router will assert a RESET signal 972 that is applied to the two clock synchronization FIFOs 518 contained in the input logic 505(sub(4)), 505(sub(5)) of the... receive symbols directly from CPUs 12A, 12B. RESET, while asserted, will hold the two clock synchronization FIFOs 518 in a temporarily non-operating reset state with the push and pull pointer... As each of the CPUs 12 receive SYNC symbols are detected by the storage and processing units of the packet receivers 96 (Figs. 5 and 6) cause the RESET signal to be asserted by the packet receivers 96 (actually, storage and processing elements 110; Fig. 6) of each CPU 12. the RESET signal is applied to the... 4)), CPUs 12 and routers 14A, 14B de-assert the RESET signals, and the clock synchronization FIFOs of the CPUs 12A, 12, and routers 14A, 14B are released from their reset... the delay, the router 14A and CPUs 12 resume pulling data from their respective clock synchronization FIFOs and resume normal operation. The clock synchronization FIFOs of the router 14A begin pulling symbols from the queue (previously set by RESET... from the CPU 12A with the T(underscore)Clk will be pushed onto the clock synchronization FIFO at, for example, queue location 0 (or whatever other location pointed to by the... 0 (or whatever other location the push pointer was set to by RESET). The clock synchronization FIFOs of the router 14A are now synchronized to accommodate whatever delay 640 may be present in one communications path, relative to the... and the CPUs 12A, 12B.

Similarly, at the same virtual time, operation of the clock synchronization FIFOs 102 of both CPUs 12A, 12B is resumed, synchronizing them to the router 14A. Also, the CPUs 12A, 12B quit sending the SLEEP command... in favor of READY symbols, and resume message packet transmission, as appropriate.

That completes the synchronization process for the router 14A. However, the process must also be performed for the router 14B. Thus, the CPU 12A returns to step... however, assuming that the CPUs 12A, 12B are operating in duplex mode, the method and apparatus used to detect and handle a possible error, resulting in divergence of the CPUs from... via a message packet destined for a peripheral device of one or the other sub-processor systems 10A, 10B. Depending upon the destination of the outgoing message packet, step 1002 will... router 14 will issue an ERROR signal to the router control logic 509, causing the process to move to step 1004 where the router 14 detecting divergence will transmit a DVRG... time outs to occur. A router detecting divergence (without also detecting any simple link error) buys itself time to check the CRC of the received message packet by waiting for the... router 14, or received, all further message packets received from the CPUs and in the process of being routed

when divergence was detected, or the DVRG symbol received, will be passed... 1010) contained in one of the configuration registers 74 (Fig. 5) of the interface unit 24 of each CPU.

Returning for the moment to step 1006, the determination of which... "local" is meant to refer to the router 14A, 14B contained in the same sub-processor system 10A, 10B as the CPU. For example, referring to Fig. 1A, router 14A is... bit mentioned above: the bit contained in one of the configuration registers 74 of interface unit 24 (Fig. 5) of each CPU. When set to a first state, that particular CPU... the other CPU. In response, the state machines (not shown) within the control and status unit 509 (Fig. 19A) changes the "favorite" bits described above.

A few examples may facilitate understanding... DVRG symbol will echo that symbol to the routers 14A, 14B, start its internal divergence process timer, and begin determination of whether to continue or terminate. Having received a TLB symbol... to diverge with no errors reported. This can happen only if software (running on the processors 20) uses known divergent data to alter state. For example, suppose each CPU 12 has... number of the CPU 12A will differ from that of the CPU 12B. If the processors use the serial number to change the sequence of instructions executed (say, by branching if the serial number comes after some value) or to modify the value contained in a processor register, the complete "state" of the CPUs 12 will differ. In such cases, the "asymmetrical... of the primary CPU simply allows one CPU, and thereby the system 10, to continue processing without software intervention.

- An error at the output of the interface unit 24 of a CPU 12 will be detected by the router 14A, 14B, depending upon... router 14A, 14B that connects to a CPU 12 will be detected by the interface unit 24 of the affected CPU. The CPU will send a TLB symbol to the faulty... possible failure and, without external intervention, and transparently to the system user, remove the failing unit (CPU 12A or 12B, or router 14A or 14B) from the system to obviate or... reintegration." The discussion will refer to the CPUs 12A, 12B, routers 14A, 14B, and maintenance processor 18A, 18B shown forming parts of the processing system 10 illustrated in Fig. 1A. In addition, discussion will refer to the processors 20a, 20b, the interface units 24a, 24b, and the memory controllers 26a, 26b (Fig. 2) of the CPUs 12A, 12B as single units, since that is the way they function.

Reintegration is used to place two CPUs in... both of the paired CPUs at virtually the same time.

The major steps in the process of changing from simplex mode operation of the one on-line CPU to duplex mode... greater detail by the flow diagrams of Figs. 33A - 33D, generally are:

1. Setup and synchronize the two CPUs (one on-line, the other off-line) and their connected routers to... the memory of the on-line CPU to the off-line CPU, maintaining a tracking process that monitors changes in the memory of the on-line CPU that have not been... and may need to be copied over to, the off-line CPU;
3. Setup and synchronize the CPUs to run a delayed (slave) duplex mode from the same instruction stream (lock... will write the predetermined registers (not shown) of the control registers 74 in the interface units 24 of CPUs 12A and 12B, to a next state (after a soft operation) in... the off-line CPU 12B.

Next, a sequence is entered (steps 1060 - 1070) that will synchronize the clock synchronization FIFOs of the CPUs 12A, 12B and routers 14A, 14B in much the same fashion... the same steps described above in connection with the discussion of Figs. 31A, 31B to synchronize the clock synchronization FIFOs. The on-line CPU 12A will send the sequence of a SLEEP symbol, self-addressed message packet, and SYNC symbol which, with the SYNC CLK signal, operates to synchronize CPUs and routers. Once so synchronized, the on-line CPU 12A then, at step 1066, sends a Soft Reset (SRST) command... of all configuration registers and control registers (e.g., configuration registers 74 of the interface units 24) cache, and the like to memory 28 of the on-line... time to have the system 10 off-line for reintegration. For that reason, the reintegration process is performed in a manner that allows the on-line CPU to continue executing user... not match that of the off-line CPU. The reason for this is that normal processing by the processor 20 of the on-line CPU can change memory content after it has been copied... when a memory location is written in the on-line CPU 12A during the reintegration process it is marked as "dirty;" second, all

copying of memory to the off-line CPU... ..may, however, limit the ability to detect two-bit errors. But, since the memory copying process will last for a only relatively short period of time, this risk is believed acceptable... ..memory location in CPU 12A is made (either an incoming I/O write, or a processor write operation). The returning data (that was copied over to the off-line CPU) would... ..controller 26 (Fig. 2) of the on-line CPU to monitor memory locations in the process of being copied over to the off-line CPU 12B. The memory controller uses a...within the block had been written by another operation (e.g., a write by the processor 20, an I/O write, etc.), that prior write operation will flag the location in... ..still must be copied over to the off-line CPU 12B.

Returning to the reintegration process, and now to Fig. 33B, the memory tracking (AtomicWrite mechanism and using ECC to mark... ..entails writing a reintegration register (not shown; one of the configuration registers 74 of interface unit 24 - Fig. 5) to cause a reintegration (REINT) signal to be asserted. The REINT signal... ..is left alone. Throughout the incremental copy operations, the normal actions of the on-line processor will mark some memory locations dirty.

Several passes of incremental copying will need to be... ..the number of successful WriteConditional operations at the end of each pass through memory, the processors 20 can determine the effect of a given pass compared to the previous pass. When the benefits drop off, the processors 20 will give up on the precopy operations. At this point the reintegration process is ready to place the two CPUs 12A, 12B into lock-step operation.

Thus, the... ..in Fig. 33C, where at step 1100, the on-line CPU 12A momentarily halts foreground processing, i.e., execution of a user application. The remaining state (e.g., configuration registers, cache, etc.) of the on-line processors 20 and its caches is then read and written to a buffer (series of memory... ..to the off-line CPU 12B, together with a "reset vector" that will direct the processor units 20 of both CPUs 12A, 12B to a reset instruction.

Next, step 1106 will quiesce... ..to ensure that the FIFOs of the routers are clear, that the FIFOs of the processor interfaces 24 are clear, and no further incoming I/O message packets are forthcoming. At... ..symbol will be received and acted upon by both CPUs 12A, 12B, to cause the processor units 20 of each CPU to jump to the location in memory 28 containing the reset... ..a subroutine that will restore the stored state of both CPUs 12A, 12B to the processor units 20, caches 22, registers, etc. The CPUs 12A, 12B will then begin executing the same... ..enabling of the ECC bit to mark dirty locations must now be disabled, since the processors are doing the same thing to the same memory. During this stage of the reintegration... ..encountered by CPU 12A.

Meanwhile, the bus error in the CPU 12A will cause the processor unit 20 to be forced into an error-handling routine to determine (1) the cause of... ..error was caused by an attempt to read a memory location marked dirty. Accordingly, the processor unit 20 will initiate (via the BTE 88 -- Fig. 5) the AtomicWrite mechanism to copy the...the SRST symbols are now received by the CPUs 12A, 12B, they will cause both processor units 20 of the CPUs to be reset to start from the same location with the... ..will periodically update, e.g., a database or audit file that is indicative of the processing of the primary CPU up to that point in time of the update. Should the... ..in error-checking redundancy to the CPU 12B, in the same manner that the individual processor units 20a, 20b of the CPU 12A provide fail-fast, fault tolerance for the CPU - when... ..cost system is applicable , as illustrated in Fig. 34. As shown in Fig. 34, a processing system 10' includes the CPU 12A and routers 14A, 14B structured as described above. The... ..and the CPUs are also the same.

Thus, the CPU 12B' comprises only a single processor unit 20' and associated support components, including the cache 22', interface unit (IU) 24', memory controller 26', and memory 28'. Thus, while the CPU 12A is structured in the manner shown in Fig. 2, with cache processor unit, interface unit, and memory control redundancies, approximately one-half of those components are needed to implement CPU... ..stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... ..i.e, parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of... inputs connected to receive the data output from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the... attached, without the assistance of any other CPU 12 in the system. Many prior parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O device becomes... rest of the system. Other prior systems provide access to I/O through pairs of processors so that should one of the processors fail, ...if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... multiprocessing system access to every peripheral, as done here, operates to extend the "primary"/"backup" process taught in the above-identified U.S. Patent No. 4,228,496. There, a multiple CPU system may have a primary process may running on one CPU, while a backup process resides in the background on another of the CPUs. Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be... e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system... CPU can access any I/O without the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time... indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Specification: ...both of the paired CPUs at virtually the same time.

The major steps in the process for changing from simplex mode operation of the one on-line CPU to duplex mode... greater detail by the flow diagrams of Figs. 33A - 33D, generally are:

1. Setup and synchronize the two CPUs (one on-line, the other off-line) and their connected routers to... the memory of the on-line CPU to the off-line CPU, maintaining a tracking process that monitors changes in the memory of the on-line CPU that have not been... and may need to be copied over to, the off-line CPU;
3. Setup and synchronize the CPUs to run a delayed (slave) duplex mode from the same instruction stream (lock... will write the predetermined registers (not shown) of the control registers 74 in the interface units 24 of CPUs 12A and 12B, to a next state (after a soft operation) in... the off-line CPU 12B.

Next, a sequence is entered (steps 1060 - 1070) that will synchronize the clock synchronization FIFOs of the CPUs 12A, 12B and routers 14A, 14B in much the same fashion... the same steps described above in connection with the discussion of Figs. 31A, 31B to synchronize the clock synchronization FIFOs. The on-line CPU 12A will send the sequence of a SLEEP symbol, self-addressed message packet, and SYNC symbol which, with the SYNC CLK signal, operates to synchronize CPUs and routers. Once so synchronized, the on-line CPU 12A then, at step 1066, sends a Soft Reset (SRST) command... of all configuration registers and control registers (e.g., configuration registers 74 of the interface units 24) cache, and the like to memory 28 of the on-line CPU, copying the... time to have the system 10 off-line for reintegration. For that reason, the reintegration process is performed in a manner that allows the on-line CPU to continue executing user... not match that of the off-line CPU. The reason for this is that normal processing by the processor 20 of the on-line CPU can change memory content after it has been copied... when a memory location is written in the on-line CPU 12A during the reintegration process it is marked ...may, however, limit the ability to detect two-bit errors. But, since the memory copying process will last for a only relatively short period of time, this risk is believed acceptable... memory location in CPU 12A is made (either an incoming I/O write, or a processor write operation). The returning data (that was copied over to the off-line CPU) would... controller 26 (Fig. 2) of the on-line CPU to monitor memory locations in the process of being copied over to the off-line CPU 12B. The memory controller uses a... within the block had been written by another operation (e.g., a write by the processor 20, an I/O write, etc.), that prior write operation will flag the location in... still must be copied over to the off-line CPU 12B.

Returning to the reintegration process, and now to Fig. 33B, the memory tracking (AtomicWrite mechanism and using ECC to mark... entails writing a reintegration register (not shown; one of the configuration registers 74 of interface unit 24 - Fig. 5) to cause a reintegration (REINT) signal to be asserted. The REINT signal... is left alone. Throughout the incremental copy operations, the normal actions of the on-line processor will mark some memory locations dirty.

Several passes of incremental copying will need to be... the number of successful WriteConditional operations at the end of each pass through memory, the processors 20 can determine the effect of a given pass compared to the previous pass. When the benefits drop off, the processors 20 will give up on the precopy operations. At this point the reintegration process is ready to place the two CPUs 12A, 12B into lock-step operation.

Thus, the... in Fig. 33C, where at step 1100, the on-line CPU 12A momentarily halts foreground processing, i.e., execution of a user application. The remaining state (e.g., configuration registers, cache, etc.) of the on-line processors 20 and its caches is then read and written to a buffer (series of memory... to the off-line CPU 12B, together with a "reset vector" that will direct the processor units 20 of both CPUs 12A, 12B to a reset instruction.

Next, step 1106 will quiesce... to ensure that the FIFOs of the routers are clear, that the FIFOs of the processor interfaces 24 are clear, and no further incoming I/O message packets are forthcoming. At... symbol will be received and acted upon by both CPUs 12A, 12B, to cause the processor units 20 of each CPU to jump to the location in memory 28 containing the reset... a subroutine that will restore the stored state of both CPUs 12A, 12B to the processor units 20, caches 22, registers, etc. The CPUs 12A, 12B will then begin executing the same... enabling of the ECC bit to mark dirty locations must now be disabled, since the processors are doing the same thing to the same memory. During this stage of the reintegration... encountered by CPU 12A.

Meanwhile, the bus error in the CPU 12A will cause the processor unit 20 to be forced into an error-handling routine to determine (1) the cause of... error was caused by an attempt to read a memory location marked dirty. Accordingly, the processor unit 20 will initiate (via the BTE 88 -- Fig. 5) the AtomicWrite mechanism to copy the... the SRST symbols are now received by the CPUs 12A, 12B, they will cause both processor units 20 of the CPUs to be reset to start from the same location with the... will periodically update, e.g., a database or audit file that is indicative of the

processing of the primary CPU up to that point in time of the update. Should the...in error-checking redundancy to the CPU 12B, in the same manner that the individual processor units 20a, 20b of the CPU 12A provide fail-fast, fault tolerance for the CPU - when... ..cost system is applicable , as illustrated in Fig. 34. As shown in Fig. 34, a processing system 10' includes the CPU 12A and routers 14A, 14B structured as described above. The... ..and the CPUs are also the same.

Thus, the CPU 12B' comprises only a single processor unit 20' and associated support components, including the cache 22', interface unit (IU) 24', memory controller 26', and memory 28'. Thus, while the CPU 12A is structured in the manner shown in Fig. 2, with cache processor unit, interface unit, and memory control redundancies, approximately one-half of those components are needed to implement CPU... ..stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... ..i.e, parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of... ..inputs connected to receive the data output from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... ..for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the... ..attached, without the assistance of any other CPU 12 in the system. Many prior parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O device becomes... ..rest of the system. Other prim systems provide access to I/O through pairs of processors so that should one of the processors fail, access to the corresponding I/O is still available through the remaining I/O... ..if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... ..multiprocessing system access to every peripheral , as done here, operates to extend the "primary"/"backup" process taught in the above-identified U.S. Patent No. 4,228,496. There, a multiple CPU system may have a primary process running on one CPU, while a backup process resides in the background on another of the CPUs. Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system... ..CPU can access any I/O without the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... ..four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... ..an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time...indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... ..When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... ..response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Claims: ...A2

1. A method of synchronizing data sent by a pair of data transmitting sources and to receive from each identical...

6/K/21 (Item 21 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Multimedia server system and method for communicating multimedia information

Country	Number	Kind	Date
---------	--------	------	------

Abstract ...A2

Type	Pub. Date	Kind	Text
Available Text	Language	Update	Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: ...A2

Field of the Invention

The present invention relates generally to server-based storage and communication systems, and, more particularly, to a multimedia server system and method for communicating multimedia information.

Background of the Invention

Advancements in communications technology and increased consumer sophistication have challenged the distributors of multimedia programming to provide the subscribing public with entertainment... services in many of the larger broadcast markets. Most pay-per-view systems permit the consumer to choose from a relatively small number of motion picture selections for home viewing, with... viewing times.

A number of on-demand video services have been developed that permit the consumer to order desired programs for home viewing through the household telephone line. For example, U... to Bell Atlantic Network Services, discloses a sophisticated video-on-demand telephone service that provides consumer ordered video programming to a plurality of households through use of a public switched telephone... system reliability. These and other related operating expenses, however, are typically passed on to the consumer.

Importantly, conventional multimedia services fail to provide media presentation control features now expected by the sophisticated consumer after enjoying more than a decade of home entertainment through the use of a video... multimedia communication system adapted to provide on-demand service to a large number of subscribing customers.

In Fig. 1, for example, there is illustrated a generalized block diagram of a conventional... over a public switched telephone network. Movies are typically stored on one or more media servers 10, each of which is multiplexed to the PSTN 16. A telephonic ordering system 14... PSTN 16, and provides a means for accepting a pay-per-view order from a customer or user 20 over the telephone. Upon verifying the account status of a user 20, the media

server 10 typically transmits the ordered movie or program to a decoder box 22 coupled to the customer's telephone line 18. The transmitted program is continuously decoded by the decoder box 22 to provide continuous presentation of the selected program on the customer's television 24. Limitations in the transmission bandwidth of the telephone lines 18, as well... ..program received from a central archive library. After establishing a telephonic link with the central server 10 over a PSTN telephone network, a selected digitized movie is downloaded in its entirety into the disk storage system incorporated into the terminal unit disclosed in the '187 patent. This and other home communication systems that employ disk storage... ..communication methodology generally results in a commercial product that is prohibitively expensive for the average consumer. Also, such systems cannot provide instantaneous viewing of a selected multimedia program immediately upon receiving the transmission of the program signals from the server 10. Moreover, VCR-type control functionality can only be provided, if at all, after downloading...control over the presentation of a selected multimedia program at a minimal cost to the consumer. There exists a further need to provide a multimedia communication system that can efficiently distribute programming to a plurality of subscribing customers without requiring complex and typically expensive server processing hardware and software at the remote communication distribution center. The present invention fulfills these and other needs.

Summary of the Invention

The present invention is a multimedia server system and a method for communicating multimedia programming to distantly situated media control systems. The multimedia server system includes a mass storage library for storing a plurality of multimedia programs. A multimedia... ..by an asynchronous transfer mode distribution switch. The custom ordered series of program segments are processed by a local media control system to provide for the sequential presentation of the program... ..pay-per-view basis;

Fig. 3 is a generalized block diagram of a novel multimedia server for communicating a synchronous, asynchronous, or combined synchronous/asynchronous series of source program segments representative... ..is a generalized block diagram of a mass storage library portion of a novel multimedia server;

Fig. 5 is an illustration of a partial series of synchronous compressed source program segments... ..source video segments contained in the first twelve segment packets transmitted by a novel multimedia server during successive transmission windows;

Fig. 11 is a generalized block diagram of a novel intelligent set-top control system adapted to communicate with a remote multimedia server to facilitate asynchronous formatting of source program segments on a multimedia DASD received from the multimedia server preferably on an on-demand, pay-per-view basis;

Fig. 12 is a depiction of...of the lower and upper disk surfaces;

Figs. 21-22 are flow charts depicting general processing steps performed by a novel multimedia server when communicating with a subscriber's set-top control system to provide on-demand transmission... ..the subscriber's set-top control system;

Fig. 23 is a flow chart depicting general processing steps performed by a novel intelligent set-top control system when communicating with a remote multimedia server to receive on-demand transmission of source program segments representative of a selected multimedia program... ..device of the set-top control system;

Figs. 24-25 are flow charts depicting general processing steps performed by a novel intelligent set-top control system when writing a custom ordered... ..novel update-in-place formatting methodology; and

Fig. 26 is a flow chart depicting general processing steps associated with effectuating a spiral-and-hold operation of a novel multimedia direct access... ..Description of the Preferred Embodiments

The present invention, as previously indicated, relates to a multimedia server system and method for communicating multimedia information over a communication channel to distantly located media... ..demand, pay-per-view basis. The present application describes the entire multimedia communication system and process for providing

multimedia program distribution from a remote multimedia server system to a plurality of distantly located set-top control systems. As such, there are... is shown a system block diagram of a multimedia communication system employing a novel multimedia server 30 configured to communicate multimedia programs to a plurality of set-top control systems 62 concurrently over a communication channel 44. In one embodiment, the multimedia server 30 transmits a video program or other visual or audio presentation as a customized series of compressed digital source program segments to a subscribing customer's set-top control system 62 on an on-demand, pay-per-view basis. The... 62 for buffering a portion or all of the multimedia program received from the multimedia server 30.

A novel DASD formatting methodology is employed to buffer the customized series of compressed... local set-top control system 62 to a subscriber's television 24, home stereo, or computer system by use of a standard household transmission line or pair of infrared transceivers. In one embodiment, the multimedia server 30 customizes the order of the source program segments in response to formatting and configuration... a significant decrease in the complexity and cost of operating and maintaining a central multimedia server system 30 adapted for distributing media-on-demand programming to a plurality of set-top... 62. By providing local control over the presentation of a multimedia program, the central multimedia server 30 need not be configured to effectuate VCR-type control functions typically desired by the subscribing customer.

Those skilled in the art will readily appreciate the significant difficulty of simultaneously servicing VCR... distribution site during the communication of user-selected programs transmitted concurrently to a plurality of customers on an on-demand, real-time basis. Providing the subscribing customer local control of a media presentation directly through the set-top control system 62 provides... significant decrease in the bandwidth of the communication channel 44 and the amount of multimedia server 30 processing overhead that would otherwise be required to service VCR-type presentation control function requests from a plurality of pay-per-view customers.

A user of the set-top control system 62 preferably communicates with the multimedia server 30 over an existing communication channel 44, such as a cable television connection, for example. It is understood that a plurality of subscribing customers can concurrently communicate with the multimedia server 30 by use of the set-top control system 62, which may be situated proximate to or remotely from a television 24 or entertainment center within the subscribing customer's home or business establishment. A communications interface preferably couples the set-top control system... for effectuating communication over the communication channel 44.

The multimedia information transmitted from the multimedia server 30 to a plurality of set-top control systems 62 is preferably transmitted in a... audio compression standard set forth audio compression specifications that are suitable for coding audio programs processed by the multimedia server 30. It is to be understood that coding standards other than... employed to facilitate communication of video, audio, and other multimedia program signals between the multimedia server 30 and a plurality of customer set-top control systems 62 without departing from the scope and spirit of the present... of explanation, the advantages and features of the disclosed media-on-demand communication method and

apparatus will be discussed generally with reference to full-motion video. Full-motion video is useful... well-suited for illustrating the advantages of the novel media-on-demand communication method and apparatus. It is to be understood that the references hereinbelow to video media are for purposes... represent limitations on the type and nature of multimedia programs and information stored on and processed by the multimedia server 30.

MULTIMEDIA SERVER

Turning now to Figs. 3-4, there is illustrated an embodiment of a novel multimedia server 30 for storing and processing a variety of multimedia programs, and for distributing selected multimedia programs concurrently to a plurality... capable of storing mass amounts of information, typically on the order of terabytes. The multimedia server 30 may include storage and distribution devices situated at a central media distribution site or... that the mass

storage library 40 may be configured with a variety of storage and processing devices covering a diverse range of technologies, and is not limited to those depicted in... a plurality of popular or frequently requested multimedia programs. In accordance with a novel media server formatting architecture and methodology disclosed hereinbelow, a DRAM storage device 37 advantageously provides for fast... representative of programming made available over local, national, and international broadcast networks. Accordingly, a subscribing customer may request from a multiplicity of pre-produced and real-time multimedia programming selections.

In... may further include other information signal stream portions. A multimedia program ordered by a subscribing customer is preferably transmitted to the customer location as a customized, multiplexed program bitstream representative of the selected multimedia program, preferably over... embodiment, the multimedia programs that are made available in the mass storage library 40 are processed through the coder 32 and index parser 33 only once, and then stored on a... single mass storage device, or, alternatively, stored across a plurality of mass storage devices. When processed by the index parser 33, each of the compressed digital video segments 48 is preferably... to one or more staging storage devices 41. A significant advantage of the novel multimedia server 30 concerns the capability of organizing source video segments 48 in a customized manner for reception by a particular customer's set-top control system 62. A plurality of staging devices 41 permits each storage device, such as digital storage device 35, to concurrently service a plurality of customer requests and organize requested multimedia program in a customized manner. The staging devices 41 may... employed to store analog multimedia information. An analog multimedia program, when requested by a subscribing customer, is preferably transferred to the coder 32, coded by the coder 32, indexed in a... 40 may be distributed amongst the various components to optimize the overhead of the multimedia server 30. Further, analog and digital multimedia programming received over a local, national, or international broadcast... be respectively directed to a coder 32 or directly to an index parser 33 for processing of real-time multimedia information.

In Fig. 5, there is shown an illustration of a... The pack layer header generally contains a pack start code, or sync code, used for synchronization purposes, and a system clock value. The system header generally contains a variety of information... functionality of a subscriber's local set-top control system 62 adapted to receive and process the customized video signal stream 54, and the manner in which a subscribing customer desires to control the presentation of a requested multimedia program.

The controller 34 preferably controls... packs in accordance with MPEG terminology, of video segments 48 concurrently to one or more customer set-top control systems 62 over the communication channel 44. It is to be understood that one or more buffer memory devices (not shown) may be employed when synchronizing the transmission of video segments 48 comprising a multiplexed signal stream between the video parser 38 and the distribution switch 42, and for synchronizing segment packet transmission between the distribution switch 42 and the communication channel 44.

It is... on the mass storage device 35 to facilitate efficient transmission of one or more pre-processed, standard customized video signal streams 54 to customer set-top control systems 62 having a predefined storage capacity and control function capability. Use of such pre-processed customized video signal streams retrieved from the mass storage device 35 obviates repetitive parsing operations... a particular set-top control system's unique configuration and presentation control functionality. Generally, the process of encoding a multimedia program requires significantly greater processing resources and a correspondingly greater processing cost as compared to decoding operations. Pre-processing or encoding multimedia programs in a manner amenable to such standardized set-top control system 62 disproportionately shifts the processing overhead to the multimedia server 30, as well as the concomitant processing costs which can be shared by the subscribing customers. It is noted that prior to transmitting a video program to a subscribing customer's set-top control system 62, the subscriber's account status is preferably verified by a billing system 36 coupled to the controller 34 of the multimedia server 30. After proper account verification is confirmed, the subscribing customer is granted authorization rights to receive multimedia programming from the multimedia server 30 preferably on a pay-per-view basis.

In Figs. 7 and 8, there are... entire video program, such as a feature-length movie or theatrical performance, for example, is processed by the coder 32 and index parser 33 into a sequential series 46 of compressed...then be transmitted in a sequential manner over the communication channel 44 to a subscribing customer's set-top control system 62. A subscribing customer's set-top control system 62 preferably includes a moderate amount of local storage, typically... 10 megabytes, for receiving the compressed sequential video signal stream 46 transmitted from the multimedia server 30. Dynamic Random Access Memory (DRAM) or a DASD may be employed to buffer the... the received compressed sequential video signal stream 46.

In accordance with this embodiment the multimedia server 30 preferably communicates concurrently with a plurality of set-top control systems 62 over a... approximately ten megabytes of internal memory, for example, the distribution switch 42 of the multimedia server 30 preferably asynchronously transmits approximately ten megabytes of multimedia program information each minute to some 600 subscribing customer locations. It is noted that a set-top control system 62 configured with a minimal amount of local memory is capable of receiving and processing the sequentially ordered compressed video signal stream 46 transmitted by the multimedia server 30, but will typically lack sufficient local memory to provide a subscriber with VCR-type... organization is not necessarily required in order to realize the advantages of the novel multimedia server 30. In the embodiment illustrated in Fig. 8, the video segments 48 processed by the video parser 38 are subdivided into one odd block, Block-A 50, and... is a function of the size of an input buffer typically provided in a subscribing customer's set-top control system 62 for the purpose of buffering packets of video segments 48 received from the multimedia server 30. The organization of each of the blocks 50 and 52 formatted as shown in... maximum packet size of ten video segments 48. As such, the input buffer of a customer's set-top control system 62 would typically be configured to store at least ten... maximum packet size of five video segments 48. As such, the input buffer of a customer's set-top control system 62 would typically be configured to store at least five... number of video segments contained in the largest video segment packet transmitted by the multimedia server 30. The additional input buffer 66 storage capacity provides for enhanced synchronization of video segments 48 being processed through the input buffer 66, and provides the multimedia server 30 with additional flexibility when asynchronously distributing video segment packets to a plurality of customer set-top control systems 62. It may be advantageously efficient, for example, for the multimedia server 30 to transmit two packets during a single transmission window to a particular set-top control system 62 to reduce server 30 processing overhead during periods of peak utilization.

Referring now to Fig. 9, there is illustrated a... developed by the inventors. These formatting equations and guidelines are preferably employed by the multimedia server 30 to optimally organize a segmented multimedia program in response to various performance and functional... set-top control system 62 adapted to receive the multimedia program transmission from the multimedia server 30.

In general, a customized video signal stream 54 preferably includes an initial asynchronous or... control system 62, and is preferably the portion of the multimedia program over which a customer has full local VCR-type presentation control. Further, as will be discussed in detail hereinbelow, the asynchronous portion of the multimedia program is concurrently buffered on the customer's set-top control system 62 while being processed for immediate display on an attached television 24 or monitor, thereby providing a subscribing customer with true on-demand viewing of a selected multimedia program. It is to be understood... system 62 adapted to receive a customized video signal stream 54 transmission from the multimedia

server 30 must generally include sufficient memory to buffer all or at least a portion of... its original temporal organization. It is important to note that cooperative operation between the multimedia server 30 and a set-top control system 62 provides for a media-on-demand communication system capable of concurrently servicing a plurality of subscribing customers, with each customer having full local VCR-type control over the presentation of a portion of the multimedia... distribution switch 42 provides for a dramatic reduction in communication channel 44 bandwidth and multimedia server 30 processing overhead in comparison to conventional video communication

systems. By transmitting each of the compressed video... data are in accordance with known Synchronous Transfer Mode (STM) methodologies.

The primary ATM information unit is the cell. ATM standards define a fixed-size cell with a length of 53... relative priority of the cell. It is noted that higher priority cells are granted preferred processing status over lower priority cells during congested intervals.

Each cell typically includes a header error.....an ATM communication network suitable for communicating a plurality of multimedia programs from a multimedia server 30 concurrently to a plurality of set-top control systems 62 preferably conforms to the... In one embodiment, the distribution architecture and method for distributing multimedia information from the multimedia server 30 to a plurality of distantly located set-top control systems 62 preferably conforms to... duration of which is preferably determined by the configuration and functional attributes of a particular customer's set-top control system 62. The customized non-sequential series of video segments 48... has an even address index, such as A2. Accordingly, an input buffer provided in a customer's set-top control system 62 would be configured to store at least two video... configured to store in excess of the minimum required capacity to provide for increased multimedia server 30 transmission flexibility and enhanced input buffer processing synchronization. In this example, an input buffer configured to store three or four video segments 48, rather than the required minimum of... an overflow buffer or transfer buffer could also be employed in cooperation with the input buffer to facilitate efficient synchronization.

By way of further example, a customized non-sequential series of video segments 48 read... would contain only four video segments 48. As such, the input buffer provided in a customer's set-top control system 62 would be configured to store at least five video... information packets unrelated to the instant multimedia program selection may also be transmitted to a customer's set-top control system 62 from the multimedia server 30. The packets containing the unrelated information, such as a message indicating that a video... vary depending on the formatting of the source program signal stream transmitted from the multimedia server 30. In general, a subscribing customer's service costs decrease as the video segment packet size transmitted by the multimedia server 30 increases. Video segment packets containing two one-segment video segments 48, for example, must be transmitted within a relatively short transmission window of approximately two seconds. The multimedia server 30 must, therefore, transmit video packets on a frequent basis. In contrast, a source multimedia... a novel intelligent set-top control system 62 adapted for communicating with a remote multimedia server 30 preferably of the type described hereinabove. In accordance with one embodiment, a relatively low... signal stream 46 comprised of sequentially ordered discrete video segments 48 transmitted from the multimedia server 30 over a communication channel 44. The set-top control system 62 preferably includes a... set-top control system 62 will generally require relatively frequent packet transmissions for the multimedia server 30, thereby resulting in higher service costs in comparison to set-top control systems employing... in accordance with a novel formatting methodology disclosed hereinbelow. An important feature afforded a subscribing customer when employing a set-top control system 62 in accordance with this embodiment concerns the... amount of available DASD 68 storage capacity generally impacts the degree to which a subscribing customer can effectuate VCR-type control over the presentation of a selected multimedia program. As illustrated... video segments 48 comprising the two-hour movie is transmitted only once from the multimedia server 30 to the subscriber's set-top control system 62. Moving outside of the presentation... compressed video segments 48. Such incidents of re-transmission preferably result in additional costs being charged to the subscriber's account.

With further reference to Fig. 11, the set-top controller 64 of the set-top control system 62 preferably communicates with a remote multimedia server 30 over a communication channel 44, and coordinates the operation of the set-top control system 62. Media-on-demand data is generally transmitted from the multimedia server 30 to the set-top control system 62 over the communication channel 44 at a... coordinate the reception, storage, and decoding of compressed video segments 48 received from the multimedia server 30, and the presentation of the decoded video segments 48 on a subscribing customer's television 76. The set-top controller 64 preferably communicates control

signals to the multimedia server 30 over a server control line or channel 78 of the communication channel 44 to initiate transmission of a... ..regulate the rate at which the compressed video signal stream is received from the multimedia server 30 over the data channel 75 to avoid an input buffer 66 overflow condition.

During... ..a control signal is preferably issued by the set-top controller 64 to the multimedia server 30 over the server control line 78 to request temporary halting of source video signal stream transmission, thus causing... ..remain stationary. The set-top controller 64 preferably issues a resume control command over the server control line 78 when requesting the multimedia server 30 to resume transmission of the source video signal stream. By way of further example, a subscribing customer may view portions of the multimedia program outside of the presentation control window 90 by... ..alerts the subscriber that satisfying the request will require additional video data from the multimedia server 30 and result in an associated charge to the subscriber's account. A subscriber may initiate transmission of the additional video data... ..signals to the input buffer 66, DASD 68, output buffer 72, decoder 74, and multimedia server 30 to regulate timing and data transmission within the set-top control system 62 respectively... ..temporarily buffering video segments being transferred into and out of the DASD 68 to enhance synchronization, and to buffer information packets and other data unrelated to the video segment 48 data prior to being... ..number is preferably used as an identification address when routing video data from the multimedia server 30 to the set-top control system 68 of the subscribing customer who placed the pay-per-view order. As discussed previously hereinabove, an ATM information cell...be situated at an outer diameter disk location, an inner diameter disk location, or an intermediate diameter disk location. It is noted that a nominal disk 108 rotation rate should be... ..the actuator performs a seek to locate a new track, it must generally decelerate and settle to a position in which it is following the centerline of the data track. Generally, a longer period of time is required for the actuator to settle at the end of a seek operation for narrower track widths, thereby increasing the overall... ..operations are performed, and, as a result, the time required for the actuator 112 to settle is no longer a significant factor that might otherwise limit the degree to which the... ..Many read errors are often imperceivable to the viewing or listening observer. Moreover, various signal processing and smoothing techniques may be employed to enhance the audio and video presentation upon the... ..Kbpi (Kilobits per inch). It is noted that more data can be stored per linear unit of track length in a spiral data track in comparison to conventional concentric tracks due...in the previously identified related U.S. Patent Application Serial No. 08/288,525 entitled "Apparatus and Method for Providing Multimedia Data."

MULTIMEDIA DASD DATA STORAGE ARCHITECTURE

Local customized control over... ..accessing of non-sequentially ordered and sequentially ordered video segments 54 received from a multimedia server 30 preferably of the type previously described. For purposes of clarity and simplicity of explanation... ..only, and do not represent limitations as to the scope of the disclosed method and apparatus.

Referring now to Figs. 15-19, it is assumed, for purposes of explanation, that the... ..window 90 for effectuating full VCR-type presentation control functions is twenty seconds, and the customer selected movie is two-hours in duration. It is noted that a typical presentation control... ..control system 62 is configured to store two discrete video segments 48. Accordingly, the multimedia server 30 transmits video segment packets containing no more than two video segments 48 to the... ..cost configuration. Such a low-cost configuration typically requires frequent packet transmissions from the multimedia server 30, thereby increasing the service costs associated with receiving multimedia programming from the multimedia server 30.

Further, it is assumed that the MPEG-1 compression standard is employed to obtain...110 and 111 is respectively shown by the direction arrows provided in Fig. 19. This

process

of sweeping over one surface of the disk 108, performing a head switch operation, and... control window 90, thus requiring transmission of additional video segment 48 information from the multimedia server 30.

PRESENTATION CONTROL WINDOW ARCHITECTURE

Still referring to Figs. 18 and 19 in detail, it...message announcing the reception of an incoming communication from a source other than the multimedia server 30, for example, may be received by the input buffer 66 and transferred to the...video segment

T_0 = Decompressed full-motion program time in seconds per video segment

P = Maximum server packet size in number of video segments based on subscriber's input buffer capacity in...
...Input buffer of set-top control system is preferably configured to store at least two server packets (P) to allow server flexibility when asynchronously transmitting video segment packets (i.e., $IBS > 2 \times P \times S_0$... video segment 48. Assuming that the maximum size of each packet transmitted by the multimedia server 30 is two segments ($P = 2$), the set-top control system's 62 input buffer...having an average size of 0.167 MB, or approximately 5 MB. For increased multimedia server 30 asynchronous transmission flexibility, the input buffer should have a storage capacity of approximately 10... a multimedia communication system. The formatting of the multimedia information received from a remote multimedia server 30 may be varied in accordance with the operational characteristics, specifications, and functions of a... multimedia program. In one illustrative embodiment, the video segment 48 information transmitted by the multimedia server 30 to the set-top control system 62 in discrete packets shown in Fig. 10... video segments (S_0), for example, multimedia program information can be efficiently transmitted from the multimedia server 30 in a format specifically tailored to the system configuration and control functionality of a subscribing customer's unique set-top control system 62.

ASYNCHRONOUS FORMATTING METHODOLOGY

Turning now to Figs. 21... depicted in chart form in Figs. 20. A subscriber preferably communicates with a remote multimedia server 30 through a novel set-top control system 62 preferably of a type discussed in... pay-per-view basis, at step 300. In this example, it is assumed that the customer is interested in selecting among various video programs, such as feature-length movies. At step... art, or a point-and-click interface similar to that commonly used when communicating with computer systems, for example.

At step 304, the subscriber preferably specifies the duration or capacity of... as a fifty minute lecture previously recorded at a local university, for example, a subscribing customer may wish to specify a fifty minute duration for the presentation control window 90 so...system 62, among others.

After selecting one or more desired multimedia programs from the multimedia server menu, the set-top controller 64, at step 308, preferably performs various internal computations to determine the nominal DASD 68 storage capacity needed to support the customer-specified presentation control window 90. The predetermined time duration (PTD) specified by the subscriber is... associated with the presentation control window 90 and the configuration and functionality of a subscribing customer's set-top control system 62 are transmitted to the multimedia server 30. The multimedia server 30 preferably includes a server controller 34 that, at step 316, reads the configuration parameters received from a subscriber's... can accommodate the subscriber-specified presentation control window 90 may instead be performed by the server controller 34 based upon the received set-top control system 62 configuration parameters.

As previously discussed, a selected multimedia program may be stored in the multimedia server 30 in either an analog format or a digital format. A selected multimedia program stored... a local, national, or international network broadcast channel 45, is typically received by the multimedia server 30 in an analog format, and may also be digitized at step 318. The digitized... segmentizing operations of steps 318 and 320 are typically not applicable to multimedia programs previously processed and stored in the multimedia server 30 in a digital format. These steps

are preferably performed only once when initially storing a multimedia program on a digital storage device 35 within the multimedia server 30.

Turning now to Fig. 22 the sequential program segments comprising the selected multimedia program are preferably arranged in a customized order at step 330. In one embodiment, the multimedia server 30 includes a video parser 38 that preferably transforms the sequential program segments into a... 62. In response to a subscriber's configuration parameters, the controller 34 of the multimedia server 30 preferably determines the ...derived the Block Indexing Coefficient, $BI = \text{modulo } D \times M$, at step 332. Further, the server controller 34, at step 334, preferably determines the length (L) of each of the segment... of rows of each customized matrix comprising a segment block (M).

At step 336, the server controller 34 preferably computes the size of the program segment packets, which is typically dependent on the size (IBS) of the input buffer 66 of a subscribing customer's set-top control system 62. The input buffer 66 of a set-top control... is transmitted to a particular set-top control system 62, is preferably computed by the server controller 34 in a manner previously discussed hereinabove. At step 340, the program segments 48... 68 formatting methodology. The set-top control system 62 formatting parameters transmitted by the multimedia server 30 are received by the subscriber's set-top control system 62 at step 350... parameters provide information preferably used by the set-top controller 64 to properly buffer and process the customized program segment packets received over the communication channel 44. For example, the segment... in turn, preferably coordinates writing and reading of the program segments received from the multimedia server 30 to and from a corresponding number (M) of storage blocks of a predetermined length... 108 surface.

At step 351, a customized program segment 48 series transmitted from the multimedia server 30 in packets over the communication channel 44 are received by the subscriber's set-top control system 62. Generally, one packet is received during each multimedia server 30 transmission window, although multiple packets may be transmitted during this time if the input... discussed, concerns the concurrent buffering and displaying of program segments 48 received from the multimedia server 30 to facilitate virtually instantaneous on-demand viewing of a selected multimedia program. As previously... 358. It is noted that, at steps 356 and 358, it may be advantageous for synchronization purposes to transfer a number of non-sequential program segments 48 received from the communication... have been transferred to the output buffer 76 in sequential order.

As mentioned previously, the process of concurrently transferring non-sequential program segments 48 to both the DASD 68, at step... 74 thus provides for virtually instantaneous presentation of a selected multimedia program on a subscribing customer's television 76. The program segments 48 contained in subsequently received Packet-3 and Packet... and, if an overflow condition is imminent, preferably transmits a control signal to the multimedia server 30 to request temporary halting of the transmission of program segment packets at step 362... periods of normal program viewing since the transmission and reception of program segment packets is synchronized by transferring packets during prescribed transmission windows. Various presentation control window 90 function modes, such... of a halt control signal from the set-top control system 62 to the multimedia server 30. The remaining program segments 48 contained in Packet-1 and stored in the input... segment A12 has been read from Location-18 during RUN-3 and displayed on the customer's television 74, and that the operations described in Figs. 24-26 are associated with... is read from Location-4, and decoded and then displayed at step 410 on the customer's television 76 in sequence with respect to the previously read and displayed program segment... is read from Location-14, and decoded and then displayed at step 430 on the customer's television 76 in sequence with respect to the previously read program segment A14. The...

Claims: ...A2

1. A server system for communicating multimedia programming to a distantly situated local media control system over a communication channel, the server system comprising:

a mass storage library for storing a plurality of multimedia programs;

organizing means... ..series to the communication channel.

2. A system as claimed in Claim 1, wherein:

the server system further comprises coding means for coding the source program segments defining the plurality of... ..with a predetermined coding standard.

3. A system as claimed in Claim 1, wherein:

the server system further comprises a server controller for communicating with the local media control system and for receiving a configuration signal... ..the communication channel.

12. A method for communicating a multimedia program from a remote multimedia server to a local media control system, the method comprising:

organizing sequentially ordered source program segments...

6/K/22 (Item 22 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Country	Number	Kind	Date
---------	--------	------	------

Abstract ...A3

Type	Pub. Date	Kind	Text
Available Text	Language	Update	Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: ...method for locally controlling multimedia programming received from a remote media-on-demand communication system server. **BACKGROUND OF THE INVENTION**

Advancements in communications technology and increased consumer sophistication have challenged the distributors of multimedia programming to provide the subscribing public with entertainment... ..services in many of the larger broadcast markets. Most pay-per-view systems permit the consumer to choose from a relatively small number of motion picture selections for home viewing, with... ..viewing times.

A number of on-demand video services have been developed that permit the consumer to order desired programs for home viewing through the household telephone line. For example, U... ..to Bell Atlantic Network Services, discloses a sophisticated video-on-demand telephone service that provides consumer ordered video programming to a plurality of households through use of a public switched telephone... ..system reliability. These and other related operating expenses, however, are typically passed on to the consumer.

Importantly, conventional multimedia services fail to provide media presentation control features now expected by the sophisticated consumer after enjoying more than a decade of home entertainment through the use of a video...

...multimedia communication system adapted to provide on-demand service to a large number of subscribing customers.

In Fig. 1, for example, there is illustrated a generalized block diagram of a conventional... over a public switched telephone network. Movies are typically stored on one or more media servers 10, each of which is multiplexed to the PSTN 16. A telephonic ordering system 14... PSTN 16, and provides a means for accepting a pay-per-view order from a customer or user 20 over the telephone. Upon verifying the account status of a user 20, the media server 10 typically transmits the ordered movie or program to a decoder box 22 coupled to the customer's telephone line 18. The transmitted program is continuously decoded by the decoder box 22 to provide continuous presentation of the selected program on the customer's television 24. Limitations in the transmission bandwidth of the telephone lines 18, as well... program received from a central archive library. After establishing a telephonic link with the central server 10 over a PSTN telephone network, a selected digitized movie is downloaded in its entirety into the disk storage system incorporated into the terminal unit disclosed in the '187 patent. This and other home communication systems that employ disk storage... communication methodology generally results in a commercial product that is prohibitively expensive for the average consumer. Also, such systems cannot provide instantaneous viewing of a selected multimedia program immediately upon receiving the transmission of the program signals from the server 10. Moreover, VCR-type control functionality can only be provided, if at all, after downloading... a need in the communications industry for a multimedia control system and method for receiving, processing, and locally controlling the presentation of a selected multimedia program transmitted from a remote media... to provide such a multimedia control system and method at a minimal cost to the consumer. The present invention fulfills these and other needs.

SUMMARY OF THE INVENTION

The present invention is an apparatus and method for effectuating local reception and processing of source program signals representative of a multimedia program received from a remote multimedia server. The multimedia server transmits a selected multimedia program as a custom ordered series of discrete, digitally compressed program... device adapted to buffer a predetermined number of compressed program segments received from a multimedia server, some of which may be non-sequentially ordered and others of which may be sequentially... formatting methodology also provides concurrent presentation and buffering of program segments received from the multimedia server for on-demand viewing of a selected multimedia program.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig... pay-per-view basis;

Fig. 3 is a generalized block diagram of a novel multimedia server for communicating a synchronous, asynchronous, or combined synchronous/asynchronous series of source program segments representative... is a generalized block diagram of a mass storage library portion of a novel multimedia server;

Fig. 5 is an illustration of a partial series of synchronous compressed source program segments... source video segments contained in the first twelve segment packets transmitted by a novel multimedia server during successive transmission windows;

Fig. 11 is a generalized block diagram of a novel intelligent set-top control system adapted to communicate with a remote multimedia server to facilitate asynchronous formatting of source program segments on a multimedia DASD received from the multimedia server preferably on an on-demand, pay-per-view basis;

Fig. 12 is a depiction of... of the lower and upper disk surfaces;

Figs. 21-22 are flow charts depicting general processing steps performed by a novel multimedia server when communicating with a subscriber's set-top control system to provide on-demand transmission... the subscriber's set-top control system;

Fig. 23 is a flow chart depicting general processing steps performed by a novel intelligent set-top control system when communicating with a remote multimedia server to receive on-demand transmission of source program segments representative of a selected multimedia program... ..device of the set-top control system;

Figs. 24-25 are flow charts depicting general processing steps performed by a novel intelligent set-top control system when writing a custom ordered... ..novel update-in-place formatting methodology; and

Fig. 26 is a flow chart depicting general processing steps associated with effectuating a spiral-and-hold operation of a novel multimedia direct access... ..the presentation of selected multimedia programs received in a customized format from a remote multimedia server, preferably on an on-demand, pay-per-view basis. The present application describes the entire multimedia communication system and process for providing multimedia program distribution from a remote multimedia server to a plurality of local set-top control systems. As such, there are described in... ..is shown a system block diagram of a multimedia communication system employing a novel multimedia server 30 configured to communicate multimedia programs to a plurality of set-top control systems 62 concurrently over a communication channel 44. In one embodiment, the multimedia server 30 transmits a video program or other visual or audio presentation as a customized series of compressed digital source program segments to a subscribing customer's set-top control system 62 on an on-demand, pay-per-view basis. The... ..62 for buffering a portion or all of the multimedia program received from the multimedia server 30.

A novel DASD formatting methodology is employed to buffer the customized series of compressed... ..local set-top control system 62 to a subscriber's television 24, home stereo, or computer system by use of a standard household transmission line or pair of infrared transceivers. In one embodiment, the multimedia server 30 customizes the order of the source program segments in response to formatting and configuration... ..a significant decrease in the complexity and cost of operating and maintaining a central multimedia server system 30 adapted for distributing media-on-demand programming to a plurality of set-top... ..62. By providing local control over the presentation of a multimedia program, the central multimedia server 30 need not be configured to effectuate VCR-type control functions typically desired by the subscribing customer.

Those skilled in the art will readily appreciate the significant difficulty of simultaneously servicing VCR... ..distribution site during the communication of user-selected programs transmitted concurrently to a plurality of customers on an on-demand, real-time basis. Providing the subscribing customer local control of a media presentation directly through the set-top control system 62 provides... ..significant decrease in the bandwidth of the communication channel 44 and the amount of multimedia server 30 processing overhead that would otherwise be required to service VCR-type presentation control function requests from a plurality of pay-per-view customers.

A user of the set-top control system 62 preferably communicates with the multimedia server 30 over an existing communication channel 44, such as a cable television connection, for example. It is understood that a plurality of subscribing customers can concurrently communicate with the multimedia server 30 by use of the set-top control system 62, which may be situated proximate to or remotely from a television 24 or entertainment center within the subscribing customer's home or business establishment. A communications interface preferably couples the set-top control system... ..for effectuating communication over the communication channel 44.

The multimedia information transmitted from the multimedia server 30 to a plurality of set-top control systems 62 is preferably transmitted in a... ..audio compression standard set forth audio compression specifications that are suitable for coding audio programs processed by the multimedia server 30. It is to be understood that coding standards other than... ..employed to facilitate communication of video, audio, and other multimedia program signals between the multimedia server 30 and a plurality of customer set-top control systems 62 without departing from the scope and spirit of the present... ..of explanation, the advantages and features of the disclosed media-on-demand communication method and

apparatus will be discussed generally with reference to full-motion video. Full-motion video is useful... well-suited for illustrating the advantages of the novel media-on-demand communication method and apparatus. It is to be understood that the references hereinbelow to video media are for purposes... represent limitations on the type and nature of multimedia programs and information stored on and processed by the multimedia server 30.

MULTIMEDIA SERVER

Turning now to Figs. 3-4, there is illustrated an embodiment of a novel multimedia server 30 for storing and processing a variety of multimedia programs, and for distributing selected multimedia programs concurrently to a plurality... capable of storing mass amounts of information, typically on the order of terabytes. The multimedia server 30 may include storage and distribution devices situated at a central media distribution site or... that the mass storage library 40 may be configured with a variety of storage and processing devices covering a diverse range of technologies, and is not limited to those depicted in... a plurality of popular or frequently requested multimedia programs. In accordance with a novel media server formatting architecture and methodology disclosed hereinbelow, a DRAM storage device 37 advantageously provides for fast... representative of programming made available over local, national, and international broadcast networks. Accordingly, a subscribing customer may request from a multiplicity of pre-produced and real-time multimedia programming selections.

In... may further include other information signal stream portions. A multimedia program ordered by a subscribing customer is preferably transmitted to the customer location as a customized, multiplexed program bitstream representative of the selected multimedia program, preferably over... embodiment, the multimedia programs that are made available in the mass storage library 40 are processed through the coder 32 and index parser 33 only once, and then stored on a... single mass storage device, or, alternatively, stored across a plurality of mass storage devices. When processed by the index parser 33, each of the compressed digital video segments 48 is preferably... to one or more staging storage devices 41. A significant advantage of the novel multimedia server 30 concerns the capability of organizing source video segments 48 in a customized manner for reception by a particular customer's set-top control system 62. A plurality of staging devices 41 permits each storage device, such as digital storage device 35, to concurrently service a plurality of customer requests and organize requested multimedia program in a customized manner. The staging devices 41 may... employed to store analog multimedia information. An analog multimedia program, when requested by a subscribing customer, is preferably transferred to the coder 32, coded by the coder 32, indexed in a... 40 may be distributed amongst the various components to optimize the overhead of the multimedia server 30. Further, analog and digital multimedia programming received over a local, national, or international broadcast... be respectively directed to a coder 32 or directly to an index parser 33 for processing of real-time multimedia information.

In Fig. 5, there is shown an illustration of a... The pack layer header generally contains a pack start code, or sync code, used for synchronization purposes, and a system clock value. The system header generally contains a variety of information... functionality of a subscriber's local set-top control system 62 adapted to receive and process the customized video signal stream 54, and the manner in which a subscribing customer desires to control the presentation of a requested multimedia program.

The controller 34 preferably controls... packs in accordance with MPEG terminology, of video segments 48 concurrently to one or more customer set-top control systems 62 over the communication channel 44. It is to be understood that one or more buffer memory devices (not shown) may be employed when synchronizing the transmission of video segments 48 comprising a multiplexed signal stream between the video parser 38 and the distribution switch 42, and for synchronizing segment packet transmission between the distribution switch 42 and the communication channel 44.

It is... on the mass storage device 35 to facilitate efficient transmission of one or more pre-processed, standard customized video signal streams 54 to customer set-top control systems 62 having a predefined storage capacity and control function capability. Use of such pre-processed customized video signal streams retrieved from the mass

storage device 35 obviates repetitive parsing operations... a particular set-top control system's unique configuration and presentation control functionality. Generally, the process of encoding a multimedia program requires significantly greater processing resources and a correspondingly greater processing cost as compared to decoding operations. Pre-processing or encoding multimedia programs in a manner amenable to such standardized set-top control system 62 disproportionately shifts the processing overhead to the multimedia server 30, as well as the concomitant processing costs which can be shared by the subscribing customers. It is noted that prior to transmitting a video program to a subscribing customer's set-top control system 62, the subscriber's account status is preferably verified by a billing system 36 coupled to the controller 34 of the multimedia server 30. After proper account verification is confirmed, the subscribing customer is granted authorization rights to receive multimedia programming from the multimedia server 30 preferably on a pay-per-view basis.

In Figs. 7 and 8, there are... entire video program, such as a feature-length movie or theatrical performance, for example, is processed by the coder 32 and index parser 33 into a sequential series 46 of compressed... then be transmitted in a sequential manner over the communication channel 44 to a subscribing customer's set-top control system 62. A subscribing customer's set-top control system 62 preferably includes a moderate amount of local storage, typically... 10 megabytes, for receiving the compressed sequential video signal stream 46 transmitted from the multimedia server 30. Dynamic Random Access Memory (DRAM) or a DASD may be employed to buffer the... the received compressed sequential video signal stream 46.

In accordance with this embodiment the multimedia server 30 preferably communicates concurrently with a plurality of set-top control systems 62 over a... approximately ten megabytes of internal memory, for example, the distribution switch 42 of the multimedia server 30 preferably asynchronously transmits approximately ten megabytes of multimedia program information each minute to some 600 subscribing customer locations. It is noted that a set-top control system 62 configured with a minimal amount of local memory is capable of receiving and processing the sequentially ordered compressed video signal stream 46 transmitted by the multimedia server 30, but will typically lack sufficient local memory to provide a subscriber with VCR-type... organization is not necessarily required in order to realize the advantages of the novel multimedia server 30. In the embodiment illustrated in Fig. 8, the video segments 48 processed by the video parser 38 are subdivided into one odd block, Block-A 50, and... is a function of the size of an input buffer typically provided in a subscribing customer's set-top control system 62 for the purpose of buffering packets of video segments 48 received from the multimedia server 30. The organization of each of the blocks 50 and 52 formatted as shown in... maximum packet size of ten video segments 48. As such, the input buffer of a customer's set-top control system 62 would typically be configured to store at least ten... maximum packet size of five video segments 48. As such, the input buffer of a customer's set-top control system 62 would typically be configured to store at least five... number of video segments contained in the largest video segment packet transmitted by the multimedia server 30. The additional input buffer 66 storage capacity provides for enhanced synchronization of video segments 48 being processed through the input buffer 66, and provides the multimedia server 30 with additional flexibility when asynchronously distributing video segment packets to a plurality of customer set-top control systems 62. It may be advantageously efficient, for example, for the multimedia server 30 to transmit two packets during a single transmission window to a particular set-top control system 62 to reduce server 30 processing overhead during periods of peak utilization.

Referring now to Fig. 9, there is illustrated a... developed by the inventors. These formatting equations and guidelines are preferably employed by the multimedia server 30 to optimally organize a segmented multimedia program in response to various performance and functional... set-top control system 62 adapted to receive the multimedia program transmission from the multimedia server 30.

In general, a customized video signal stream 54 preferably includes an initial asynchronous or... control system 62, and is preferably the portion of the multimedia program over which a customer has full local VCR-type presentation control. Further, as will be discussed in detail hereinbelow, the asynchronous portion of the multimedia program is concurrently buffered on the customer's set-top control system 62 while being processed for immediate display on

an attached television 24 or monitor, thereby providing a subscribing customer with true on-demand viewing of a selected multimedia program. It is to be understood... ..system 62 adapted to receive a customized video signal stream 54 transmission from the multimedia

server 30 must generally include sufficient memory to buffer all or at least a portion of... ..its original temporal organization. It is important to note that cooperative operation between the multimedia server 30 and a set-top control system 62 provides for a media-on-demand communication system capable of concurrently servicing a plurality of subscribing customers, with each customer having full local VCR-type control over the presentation of a portion of the multimedia... ..distribution switch 42 provides for a dramatic reduction in communication channel 44 bandwidth and multimedia server 30 processing overhead in comparison to conventional video communication systems. By transmitting each of the compressed video...data are in accordance with known Synchronous Transfer Mode (STM) methodologies.

The primary ATM information unit is the cell. ATM standards define a fixed-size cell with a length of 53... ..relative priority of the cell. It is noted that higher priority cells are granted preferred processing status over lower priority cells during congested intervals.

Each cell typically includes a header error... ..an ATM communication network suitable for communicating a plurality of multimedia programs from a multimedia server 30 concurrently to a plurality of set-top control systems 62 preferably conforms to the... ..In one embodiment, the distribution architecture and method for distributing multimedia information from the multimedia server 30 to a plurality of distantly located set-top control systems 62 preferably conforms to... ..duration of which is preferably determined by the configuration and functional attributes of a particular customer's set-top control system 62. The customized non-sequential series of video segments 48... ..has an even address index, such as A2. Accordingly, an input buffer provided in a customer's set-top control system 62 would be configured to store at least two video... ..configured to store in excess of the minimum required capacity to provide for increased multimedia server 30 transmission flexibility and enhanced input buffer processing synchronization. In this example, an input buffer configured to store three or four video segments 48, rather than the required minimum of... ..an overflow buffer or transfer buffer could also be employed in cooperation with the input buffer to facilitate efficient synchronization.

By way of further example, a customized non-sequential series of video segments 48 read... ..would contain only four video segments 48. As such, the input buffer provided in a customer's set-top control system 62 would be configured to store at least five video... ..information packets unrelated to the instant multimedia program selection may also be transmitted to a customer's set-top control system 62 from the multimedia server 30. The packets containing the unrelated information, such as a message indicating that a video... ..vary depending on the formatting of the source program signal stream transmitted from the multimedia server 30. In general, a subscribing customer's service costs decrease as the video segment packet size transmitted by the multimedia server 30 increases. Video segment packets containing two one-segment video segments 48, for example, must be transmitted within a relatively short transmission window of approximately two seconds. The multimedia server 30 must, therefore, transmit video packets on a frequent basis. In contrast, a source multimedia... ..a novel intelligent set-top control system 62 adapted for communicating with a remote multimedia server 30 preferably of the type described hereinabove. In accordance with one embodiment, a relatively low... ..signal stream 46 comprised of sequentially ordered discrete video segments 48 transmitted from the multimedia server 30 over a communication channel 44. The set-top control system 62 preferably includes a... ..set-top control system 62 will generally require relatively frequent packet transmissions for the multimedia server 30, thereby resulting in higher service costs in comparison to set-top control systems employing... ..in accordance with a novel formatting methodology disclosed hereinbelow. An important feature afforded a subscribing customer when employing a set-top control system 62 in accordance with this embodiment concerns the... ..amount of available DASD 68 storage capacity generally impacts the degree to which a subscribing customer can effectuate VCR-type control over the presentation of a selected multimedia

program. As illustrated... video segments 48 comprising the two-hour movie is transmitted only once from the multimedia server 30 to the subscriber's set-top control system 62. Moving outside of the presentation... compressed video segments 48. Such incidents of re-transmission preferably result in additional costs being charged to the subscriber's account.

With further reference to Fig. 11, the set-top controller 64 of the set-top control system 62 preferably communicates with a remote multimedia server 30 over a communication channel 44, and coordinates the operation of the set-top control system 62. Media-on-demand data is generally transmitted from the multimedia server 30 to the set-top control system 62 over the communication channel 44 at a... coordinate the reception, storage, and decoding of compressed video segments 48 received from the multimedia server 30, and the presentation of the decoded video segments 48 on a subscribing customer's television 76. The set-top controller 64 preferably communicates control signals to the multimedia server 30 over a server control line or channel 78 of the communication channel 44 to initiate transmission of a... regulate the rate at which the compressed video signal stream is received from the multimedia server 30 over the data channel 75 to avoid an input buffer 66 overflow condition.

During... a control signal is preferably issued by the set-top controller 64 to the multimedia server 30 over the server control line 78 to request temporary halting of source video signal stream transmission, thus causing... remain stationary. The set-top controller 64 preferably issues a resume control command over the server control line 78 when requesting the multimedia server 30 to resume transmission of the source video signal stream. By way of further example, a subscribing customer may view portions of the multimedia program outside of the presentation control window 90 by... alerts the subscriber that satisfying the request will require additional video data from the multimedia server 30 and result in an associated charge to the subscriber's account. A subscriber may initiate transmission of the additional video data... signals to the input buffer 66, DASD 68, output buffer 72, decoder 74, and multimedia server 30 to regulate timing and data transmission within the set-top control system 62 respectively... temporarily buffering video segments being transferred into and out of the DASD 68 to enhance synchronization, and to buffer information packets and other data unrelated to the video segment 48 data prior to being... number is preferably used as an identification address when routing video data from the multimedia server 30 to the set-top control system 68 of the subscribing customer who placed the pay-per-view order. As discussed previously hereinabove, an ATM information cell... be situated at an outer diameter disk location, an inner diameter disk location, or an intermediate diameter disk location. It is noted that a nominal disk 108 rotation rate should be... the actuator performs a seek to locate a new track, it must generally decelerate and settle to a position in which it is following the centerline of the data track. Generally, a longer period of time is required for the actuator to settle at the end of a seek operation for narrower track widths, thereby increasing the overall... operations are performed, and, as a result, the time required for the actuator 112 to settle is no longer a significant factor that might otherwise limit the degree to which the... Many read errors are often imperceptible to the viewing or listening observer. Moreover, various signal processing and smoothing techniques may be employed to enhance the audio and video presentation upon the... Kbpi (Kilobits per inch). It is noted that more data can be stored per linear unit of track length in a spiral data track in comparison to conventional concentric tracks due... in the previously identified related U.S. Patent Application Serial No. 08/288,525 entitled "Apparatus and Method for Providing Multimedia Data."

MULTIMEDIA DASD DATA STORAGE ARCHITECTURE

Local customized control over... accessing of non-sequentially ordered and sequentially ordered video segments 54 received from a multimedia server 30 preferably of the type previously described. For purposes of clarity and simplicity of explanation... only, and do not represent limitations as to the scope of the disclosed method and apparatus.

Referring now to Figs. 15-19, it is assumed, for purposes of explanation, that the... window 90 for effectuating full VCR-type presentation control functions is twenty seconds, and the customer selected movie is two-hours in duration. It is noted that a typical presentation control... control system 62 is configured to store two discrete video

segments 48. Accordingly, the multimedia server 30 transmits video segment packets containing no more than two video segments 48 to the... ..cost configuration. Such a low-cost configuration typically requires frequent packet transmissions from the multimedia server 30, thereby increasing the service costs associated with receiving multimedia programming from the multimedia server 30.

Further, it is assumed that the MPEG-1 compression standard is employed to obtain... ..110 and 111 is respectively shown by the direction arrows provided in Fig. 19. This

process of sweeping over one surface of the disk 108, performing a head switch operation, and... ..control window 90, thus requiring transmission of additional video segment 48 information from the multimedia server 30.

PRESENTATION CONTROL WINDOW ARCHITECTURE

Still referring to Figs. 18 and 19 in detail, it... ..message annunciating the reception of an incoming communication from a source other than the multimedia server 30, for example, may be received by the input buffer 66 and transferred to the...video segment

TO = Decompressed full-motion program time in seconds per video segment

P = Maximum server packet size in number of video segments based on subscriber's input buffer capacity in... ..Input buffer of set-top control system is preferably configured to store at least two server packets (P) to allow server flexibility when asynchronously transmitting video segment packets (i.e., $IBS > 2 \times P \times S0$video segment 48. Assuming that the maximum size of each packet transmitted by the multimedia server 30 is two segments ($P = 2$), the set-top control system's 62 input buffer... ..having an average size of 0.167 MB, or approximately 5 MB. For increased multimedia server 30 asynchronous transmission flexibility, the input buffer should have a storage capacity of approximately 10... ..a multimedia communication system. The formatting of the multimedia information received from a remote multimedia server 30 may be varied in accordance with the operational characteristics, specifications, and functions of a... ..multimedia program. In one illustrative embodiment, the video segment 48 information transmitted by the multimedia server 30 to the set-top control system 62 in discrete packets shown in Fig. 10... ..video segments (S0), for example, multimedia program information can be efficiently transmitted from the multimedia server 30 in a format specifically tailored to the system configuration and control functionality of a subscribing customer's unique set-top control system 62.

ASYNCHRONOUS FORMATTING METHODOLOGY

Turning now to Figs. 21... ..depicted in chart form in Figs. 20. A subscriber preferably communicates with a remote multimedia server 30 through a novel set-top control system 62 preferably of a type discussed in... ..pay-per-view basis, at step 300. In this example, it is assumed that the customer is interested in selecting among various video programs, such as feature-length movies. At step... ..art, or a point-and-click interface similar to that commonly used when communicating with computer systems, for example.

At step 304, the subscriber preferably specifies the duration or capacity of... ..as a fifty minute lecture previously recorded at a local university, for example, a subscribing customer may wish to specify a fifty minute duration for the presentation control window 90 so... ..system 62, among others.

After selecting one or more desired multimedia programs from the multimedia server menu, the set-top controller 64, at step 308, preferably performs various internal computations to determine the nominal DASD 68 storage capacity needed to support the customer-specified presentation control window 90. The predetermined time duration (PTD) specified by the subscriber is... ..associated with the presentation control window 90 and the configuration and functionality of a subscribing customer's set-top control system 62 are transmitted to the multimedia server 30. The multimedia server 30 preferably includes a server controller 34 that, at step 316, reads the configuration parameters received from a subscriber's... ..can accommodate the subscriber-specified presentation control window

90 may instead be performed by the server controller 34 based upon the received set-top control system 62 configuration parameters.

As previously discussed, a selected multimedia program may be stored in the multimedia server 30 in either an analog format or a digital format. A selected multimedia program stored... a local, national, or international network broadcast channel 45, is typically received by the multimedia server 30 in an analog format, and may also be digitized at step 318. The digitized... segmentizing operations of steps 318 and 320 are typically not applicable to multimedia programs previously processed and stored in the multimedia server 30 in a digital format. These steps are preferably performed only once when initially storing a multimedia program on a digital storage device 35 within the multimedia server 30.

Turning now to Fig. 22 the sequential program segments comprising the selected multimedia program are preferably arranged in a customized order at step 330. In one embodiment, the multimedia server 30 includes a video parser 38 that preferably transforms the sequential program segments into a... 62. In response to a subscriber's configuration parameters, the controller 34 of the multimedia server 30 preferably determines the number of segment blocks (M) per disk surface (D), from which... derived the Block Indexing Coefficient, $BI = \text{modulo } D \times M$, at step 332. Further, the server controller 34, at step 334, preferably determines the length (L) of each of the segment... of rows of each customized matrix comprising a segment block (M).

At step 336, the server controller 34 preferably computes the size of the program segment packets, which is typically dependent on the size (IBS) of the input buffer 66 of a subscribing customer's set-top control system 62. The input buffer 66 of a set-top control... is transmitted to a particular set-top control system 62, is preferably computed by the server controller 34 in a manner previously discussed hereinabove. At step 340, the program segments 48... 68 formatting methodology. The set-top control system 62 formatting parameters transmitted by the multimedia server 30 are received by the subscriber's set-top control system 62 at step 350... parameters provide information preferably used by the set-top controller 64 to properly buffer and process the customized program segment packets received over the communication channel 44. For example, the segment... in turn, preferably coordinates writing and reading of the program segments received from the multimedia server 30 to and from a corresponding number (M) of storage blocks of a predetermined length... 108 surface.

At step 351, a customized program segment 48 series transmitted from the multimedia server 30 in packets over the communication channel 44 are received by the subscriber's set-top control system 62. Generally, one packet is received during each multimedia server 30 transmission window, although multiple packets may be transmitted during this time if the input... discussed, concerns the concurrent buffering and displaying of program segments 48 received from the multimedia server 30 to facilitate virtually instantaneous on-demand viewing of a selected multimedia program. As previously... 358. It is noted that, at steps 356 and 358, it may be advantageous for synchronization purposes to transfer a number of non-sequential program segments 48 received from the communication... have been transferred to the output buffer 76 in sequential order.

As mentioned previously, the process of concurrently transferring non-sequential program segments 48 to both the DASD 68, at step... 74 thus provides for virtually instantaneous presentation of a selected multimedia program on a subscribing customer's television 76. The program segments 48 contained in subsequently received Packet-3 and Packet... and, if an overflow condition is imminent, preferably transmits a control signal to the multimedia server 30 to request temporary halting of the transmission of program segment packets at step 362... periods of normal program viewing since the transmission and reception of program segment packets is synchronized by transferring packets during prescribed transmission windows. Various presentation control window 90 function modes, such... of a halt control signal from the set-top control system 62 to the multimedia server 30. The remaining program segments 48 contained in Packet-1 and stored in the input... segment A12 has been read from Location-18 during RUN-3 and displayed on the customer's television 74, and that the operations described in Figs. 24-26 are associated with... is read from Location-4, and decoded and then displayed at step 410 on the customer's television 76 in sequence with respect to the previously read and displayed program segment... is read from Location-14, and

decoded and then displayed at step 430 on the customer's television 76 in sequence with respect to the previously read program segment A14. The...

Claims: ...system, coupled to a display and a communication channel, for communicating with a remote multimedia server having means for transmitting source program segments representative of a multimedia program as a custom... ...ordered local program segments.

3. A system as claimed in Claim 1, wherein the multimedia server transmits the source program segments to the communication channel arranged in packets, wherein each of... ...from the communication channel.

4. A system as claimed in Claim 1, wherein:

the multimedia server further comprises coding means for coding the source program segments defining the multimedia program as... ...claimed in Claim 1, wherein the controller communicates a control signal to the remote multimedia server to halt transmission of the source program segments to the multimedia control system to preclude... ...method for locally controlling the presentation of a multimedia program received from a remote multimedia server as a custom ordered series of coded source program segments arranged in packets, each of...

6/K/23 (Item 23 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Country	Number	Kind	Date
---------	--------	------	------

Abstract ...the direct access storage device are disclosed. A multimedia program is transmitted from a multimedia server as a custom ordered series of discrete program segments and received by the multimedia direct... ...control system for buffering a predetermined number of compressed program segments received from the multimedia server, some of which may be non-sequentially ordered and others of which may be sequentially... ...formatting methodology also provides concurrent presentation and buffering of program segments received from the multimedia server for on-demand viewing of a selected multimedia program.

Type	Pub. Date	Kind	Text
Available Text	Language	Update	Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: ...method for storing multimedia information. BACKGROUND OF THE INVENTION

Advancements in communications technology and increased consumer sophistication have challenged the distributors of multimedia programming to provide the subscribing public with entertainment... ...services in many of the larger broadcast markets. Most pay-per-view systems permit the consumer to choose from a relatively small number of motion picture selections for home viewing, with... ...viewing times.

A number of on-demand video services have been developed that permit the consumer to order desired programs for home viewing through the household telephone line. For example, U... ...to Bell Atlantic Network Services, discloses a sophisticated video-on-demand telephone service that provides consumer ordered video programming to a

plurality of households through use of a public switched telephone... ..system reliability. These and other related operating expenses, however, are typically passed on to the consumer .

Importantly, conventional multimedia services fail to provide media presentation control features now expected by the sophisticated consumer after enjoying more than a decade of home entertainment through the use of a video... ..multimedia communication system adapted to provide on-demand service to a large number of subscribing customers.

In Fig. 1, for example, there is illustrated a generalized block diagram of a conventional... ..over a public switched telephone network. Movies are typically stored on one or more media servers 10, each of which is multiplexed to the PSTN 16. A telephonic ordering system 14... ..PSTN 16, and provides a means for accepting a pay-per-view order from a customer or user 20 over the telephone. Upon verifying the account status of a user 20, the media server 10 typically transmits the ordered movie or program to a decoder box 22 coupled to the customer's telephone line 18. The transmitted program is continuously decoded by the decoder box 22 to provide continuous presentation of the selected program on the customer's television 24. Limitations in the transmission bandwidth of the telephone lines 18, as well... ..program received from a central archive library. After establishing a telephonic link with the central server 10 over a PSTN telephone network, a selected digitized movie is downloaded in its entirety into the disk storage system incorporated into the terminal unit disclosed in the '187 patent. This and other home communication systems that employ disk storage... ..communication methodology generally results in a commercial product that is prohibitively expensive for the average consumer . Also, such systems cannot provide instantaneous viewing of a selected multimedia program immediately upon receiving the transmission of the program signals from the server 10. Moreover, VCR-type control functionality can only be provided, if at all, after downloading... ..access storage device adapted to store multimedia information received from a media-on-demand communication server system, and a method for efficiently formatting multimedia information on one or more data storage... ..control over the presentation of a selected multimedia program at a minimal cost to the consumer . The present invention fulfills these and other needs.

SUMMARY OF THE INVENTION

The present invention... ..and from the direct access storage device. A multimedia program is transmitted from a multimedia server as a custom ordered series of discrete, digitally compressed program segments and received by the... ..control system for buffering a predetermined number of compressed program segments received from the multimedia server, some of which may be non-sequentially ordered and others of which may be sequentially... ..formatting methodology also provides concurrent presentation and buffering of program segments received from the multimedia server for on-demand viewing of a selected multimedia program.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig... ..pay-per-view basis;

Fig. 3 is a generalized block diagram of a novel multimedia server for communicating a synchronous, asynchronous, or combined synchronous/asynchronous series of source program segments representative... ..is a generalized block diagram of a mass storage library portion of a novel multimedia server;

Fig. 5 is an illustration of a partial series of synchronous compressed source program segments... ..source video segments contained in the first twelve segment packets transmitted by a novel multimedia server during successive transmission windows;

Fig. 11 is a generalized block diagram of a novel intelligent set-top control system adapted to communicate with a remote multimedia server to facilitate asynchronous formatting of source program segments on a multimedia DASD received from the multimedia server preferably on an on-demand, pay-per-view basis;

Fig. 12 is a depiction of... ..of the lower and upper disk surfaces;

Figs. 21-22 are flow charts depicting general processing steps performed by a novel multimedia server when communicating with a subscriber's set-top control system to provide on-demand transmission... the subscriber's set-top control system;

Fig. 23 is a flow chart depicting general processing steps performed by a novel intelligent set-top control system when communicating with a remote multimedia server to receive on-demand transmission of source program segments representative of a selected multimedia program... device of the set-top control system;

Figs. 24-25 are flow charts depicting general processing steps performed by a novel intelligent set-top control system when writing a custom ordered... novel update-in-place formatting methodology; and

Fig. 26 is a flow chart depicting general processing steps associated with effectuating a spiral-and-hold operation of a novel multimedia direct access... the presentation of selected multimedia programs received in a customized format from a remote multimedia server, preferably on an on-demand, pay-per-view basis. The present application describes the entire multimedia communication system and process for providing multimedia program distribution from a remote multimedia server to a plurality of local set-top control systems which preferably include multimedia direct access... is shown a system block diagram of a multimedia communication system employing a novel multimedia server 30 configured to communicate multimedia programs to a plurality of set-top control systems 62 concurrently over a communication channel 44. In one embodiment, the multimedia server 30 transmits a video program or other visual or audio presentation as a customized series of compressed digital source program segments to a subscribing customer's set-top control system 62 on an on-demand, pay-per-view basis. The... 62 for buffering a portion or all of the multimedia program received from the multimedia server 30.

A novel DASD formatting methodology is employed to buffer the customized series of compressed... local set-top control system 62 to a subscriber's television 24, home stereo, or computer system by use of a standard household transmission line or pair of infrared transceivers. In one embodiment, the multimedia server 30 customizes the order of the source program segments in response to formatting and configuration... a significant decrease in the complexity and cost of operating and maintaining a central multimedia server system 30 adapted for distributing media-on-demand programming to a plurality of set-top... 62. By providing local control over the presentation of a multimedia program, the central multimedia server 30 need not be configured to effectuate VCR-type control functions typically desired by the subscribing customer.

Those skilled in the art will readily appreciate the significant difficulty of simultaneously servicing VCR... distribution site during the communication of user-selected programs transmitted concurrently to a plurality of customers on an on-demand, real-time basis. Providing the subscribing customer local control of a media presentation directly through the set-top control system 62 provides... significant decrease in the bandwidth of the communication channel 44 and the amount of multimedia server 30 processing overhead that would otherwise be required to service VCR-type presentation control function requests from a plurality of pay-per-view customers.

A user of the set-top control system 62 preferably communicates with the multimedia server 30 over an existing communication channel 44, such as a cable television connection, for example. It is understood that a plurality of subscribing customers can concurrently communicate with the multimedia server 30 by use of the set-top control system 62, which may be situated proximate to or remotely from a television 24 or entertainment center within the subscribing customer's home or business establishment. A communications interface preferably couples the set-top control system... for effectuating communication over the communication channel 44.

The multimedia information transmitted from the multimedia server 30 to a plurality of set-top control systems 62 is preferably transmitted in a... audio compression standard set forth audio compression specifications that are suitable for coding audio programs processed by the multimedia sever... employed to facilitate communication of video, audio, and other multimedia program signals between the multimedia server 30 and a plurality of customer set-top control systems 62 without departing from the scope and spirit of the present... of explanation, the advantages and features of the disclosed media-on-demand communication method and apparatus will be discussed

generally with reference to full-motion video. Full-motion video is useful... well-suited for illustrating the advantages of the novel media-on-demand communication method and

apparatus. It is to be understood that the references hereinbelow to video media are for purposes... represent limitations on the type and nature of multimedia programs and information stored on and processed by the multimedia server 30.

MULTIMEDIA SERVER

Turning now to Figs. 3-4, there is illustrated an embodiment of a novel multimedia server 30 for storing and processing a variety of multimedia programs, and for distributing selected multimedia programs concurrently to a plurality... capable of storing mass amounts of information, typically on the order of terabytes. The multimedia server 30 may include storage and distribution devices situated at a central media distribution site or... that the mass storage library 40 may be configured with a variety of storage and processing devices covering a diverse range of technologies, and is not limited to those depicted in... a plurality of popular or frequently requested multimedia programs. In accordance with a novel media server formatting architecture and methodology disclosed hereinbelow, a DRAM storage device 37 advantageously provides for fast... representative of programming made available over local, national, and international broadcast networks. Accordingly, a subscribing customer may request from a multiplicity of pre-produced and real-time multimedia programming selections.

In... may further include other information signal stream portions. A multimedia program ordered by a subscribing customer is preferably transmitted to the customer location as a customized, multiplexed program bitstream representative of the selected multimedia program, preferably over... embodiment, the multimedia programs that are made available in the mass storage library 40 are processed through the coder 32 and index parser 33 only once, and then stored on a... single mass storage device, or, alternatively, stored across a plurality of mass storage devices. When processed by the index parser 33, each of the compressed digital video segments 48 is preferably... to one or more staging storage devices 41. A significant advantage of the novel multimedia server 30 concerns the capability of organizing source video segments 48 in a customized manner for reception by a particular customer's set-top control system 62. A plurality of staging devices 41 permits each storage device, such as digital storage device 35, to concurrently service a plurality of customer requests and organize requested multimedia program in a customized manner. The staging devices 41 may... employed to store analog multimedia information. An analog multimedia program, when requested by a subscribing customer, is preferably transferred to the coder 32, coded by the coder 32, indexed in a... 40 may be distributed amongst the various components to optimize the overhead of the multimedia server 30. Further, analog and digital multimedia programming received over a local, national, or international broadcast... be respectively directed to a coder 32 or directly to an index parser 33 for processing of real-time multimedia information.

In Fig. 5, there is shown an illustration of a... The pack layer header generally contains a pack start code, or sync code, used for synchronization purposes, and a system clock value. The system header generally contains a variety of information... functionality of a subscriber's local set-top control system 62 adapted to receive and process the customized video signal stream 54, and the manner in which a subscribing customer desires to control the presentation of a requested multimedia program.

The controller 34 preferably controls... packs in accordance with MPEG terminology, of video segments 48 concurrently to one or more customer set-top control systems 62 over the communication channel 44. It is to be understood that one or more buffer memory devices (not shown) may be employed when synchronizing the transmission of video segments 48 comprising a multiplexed signal stream between the video parser 38 and the distribution switch 42, and for synchronizing segment packet transmission between the distribution switch 42 and the communication channel 44.

It is... ..on the mass storage device 35 to facilitate efficient transmission of one or more pre-processed, standard customized video signal streams 54 to customer set-top control systems 62 having a predefined storage capacity and control function capability. Use of such pre-processed customized video signal streams retrieved from the mass storage device 35 obviates repetitive parsing operations... ..a particular set-top control system's unique configuration and presentation control functionality. Generally, the process of encoding a multimedia program requires significantly greater processing resources and a correspondingly greater processing cost as compared to decoding operations. Pre-processing or encoding multimedia programs in a manner amenable to such standardized set-top control system 62 disproportionately shifts the processing overhead to the multimedia server 30, as well as the concomitant processing costs which can be shared by the subscribing customers. It is noted that prior to transmitting a video program to a subscribing customer's set-top control system 62, the subscriber's account status is preferably verified by a billing system 36 coupled to the controller 34 of the multimedia server 30. After proper account verification is confirmed, the subscribing customer is granted authorization rights to receive multimedia programming from the multimedia server 30 preferably on a pay-per-view basis.

In Figs. 7 and 8, there are... ..entire video program, such as a feature-length movie or theatrical performance, for example, is processed by the coder 32 and index parser 33 into a sequential series 46 of compressed... ..then be transmitted in a sequential manner over the communication channel 44 to a subscribing customer's set-top control system 62. A subscribing customer's set-top control system 62 preferably includes a moderate amount of local storage, typically... ..10 megabytes, for receiving the compressed sequential video signal stream 46 transmitted from the multimedia server 30. Dynamic Random Access Memory (DRAM) or a DASD may be employed to buffer the... ..the received compressed sequential video signal stream 46.

In accordance with this embodiment the multimedia server 30 preferably communicates concurrently with a plurality of set-top control systems 62 over a... ..approximately ten megabytes of internal memory, for example, the distribution switch 42 of the multimedia server 30 preferably asynchronously transmits approximately ten megabytes of multimedia program information each minute to some 600 subscribing customer locations. It is noted that a set-top control system 62 configured with a minimal amount of local memory is capable of receiving and processing the sequentially ordered compressed video signal stream 46 transmitted by the multimedia server 30, but will typically lack sufficient local memory to provide a subscriber with VCR-type... ..organization is not necessarily required in order to realize the advantages of the novel multimedia server 30. In the embodiment illustrated in Fig. 8, the video segments 48 processed by the video parser 38 are subdivided into one odd block, Block-A 50, and... ..is a function of the size of an input buffer typically provided in a subscribing customer's set-top control system 62 for the purpose of buffering packets of video segments 48 received from the multimedia server 30. The organization of each of the blocks 50 and 52 formatted as shown in... ..maximum packet size of ten video segments 48. As such, the input buffer of a customer's set-top control system 62 would typically be configured to store at least ten... ..maximum packet size of five video segments 48. As such, the input buffer of a customer's set-top control system 62 would typically be configured to store at least five... ..number of video segments contained in the largest video segment packet transmitted by the multimedia server 30. The additional input buffer 66 storage capacity provides for enhanced synchronization of video segments 48 being processed through the input buffer 66, and provides the multimedia server 30 with additional flexibility when asynchronously distributing video segment packets to a plurality of customer set-top control systems 62. It may be advantageously efficient, for example, for the multimedia server 30 to transmit two packets during a single transmission window to a particular set-top control system 62 to reduce server 30 processing overhead during periods of peak utilization.

Referring now to Fig. 9, there is illustrated a... ..developed by the inventors. These formatting equations and guidelines are preferably employed by the multimedia server 30 to optimally organize a segmented multimedia program in response to various performance and functional... ..set-top control system 62 adapted to receive the multimedia program transmission from the multimedia server 30.

In general, a customized video signal stream 54 preferably includes an initial asynchronous or...control system 62, and is preferably the portion of the multimedia program over which a customer has full local VCR-type presentation control. Further, as will be discussed in detail hereinbelow, the asynchronous portion of the multimedia program is concurrently buffered on the customer's set-top control system 62 while being processed for immediate display on an attached television 24 or monitor, thereby providing a subscribing customer with true on-demand viewing of a selected multimedia program. It is to be understood... ..system 62 adapted to receive a customized video signal stream 54 transmission from the multimedia server 30 must generally include sufficient memory to buffer all or at least a portion of... ..its original temporal organization. It is important to note that cooperative operation between the multimedia

server 30 and a set-top control system 62 provides for a media-on-demand communication system capable of concurrently servicing a plurality of subscribing customers, with each customer having full local VCR-type control over the presentation of a portion of the multimedia... ..distribution switch 42 provides for a dramatic reduction in communication channel 44 bandwidth and multimedia server 30 processing overhead in comparison to conventional video communication systems. By transmitting each of the compressed video... ..data are in accordance with known Synchronous Transfer Mode (STM) methodologies.

The primary ATM information unit is the cell. ATM standards define a fixed-size cell with a length of 53... ..relative priority of the cell. It is noted that higher priority cells are granted preferred processing status over lower priority cells during congested intervals.

Each cell typically includes a header error... ..an ATM communication network suitable for communicating a plurality of multimedia programs from a multimedia server 30 concurrently to a plurality of set-top control systems 62 preferably conforms to the... ..In one embodiment, the distribution architecture and method for distributing multimedia information from the multimedia server 30 to a plurality of distantly located set-top control systems 62 preferably conforms to... ..duration of which is preferably determined by the configuration and functional attributes of a particular customer's set-top control system 62. The customized non-sequential series of video segments 48... ..has an even address index, such as A2. Accordingly, an input buffer provided in a customer's set-top control system 62 would be configured to store at least two video... ..configured to store in excess of the minimum required capacity to provide for increased multimedia server 30 transmission flexibility and enhanced input buffer processing synchronization. In this example, an input buffer configured to store three or four video segments 48, rather than the required minimum of... ..an overflow buffer or transfer buffer could also be employed in cooperation with the input buffer to facilitate efficient synchronization.

By way of further example, a customized non-sequential series of video segments 48 read... ..would contain only four video segments 48. As such, the input buffer provided in a customer's set-top control system 62 would be configured to store at least five video... ..information packets unrelated to the instant multimedia program selection may also be transmitted to a customer's set-top control system 62 from the multimedia server 30. The packets containing the unrelated information, such as a message indicating that a video... ..vary depending on the formatting of the source program signal stream transmitted from the multimedia server 30. In general, a subscribing customer's service costs decrease as the video segment packet size transmitted by the multimedia server 30 increases. Video segment packets containing two one-segment video segments 48, for example, must be transmitted within a relatively short transmission window of approximately two seconds. The multimedia server 30 must, therefore, transmit video packets on a frequent basis. In contrast, a source multimedia... ..a novel intelligent set-top control system 62 adapted for communicating with a remote multimedia server 30 preferably of the type described hereinabove. In accordance with one embodiment, a relatively low... ..signal stream 46 comprised of sequentially ordered discrete video segments 48 transmitted from the multimedia server 30 over a communication channel 44. The set-top control system 62 preferably includes a... ..set-top control system 62 will generally require relatively frequent packet transmissions for the multimedia server 30, thereby resulting in higher service costs in comparison

to set-top control systems employing... ..in accordance with a novel formatting methodology disclosed hereinbelow. An important feature afforded a subscribing customer when employing a set-top control system 62 in accordance with this embodiment concerns the... ..amount of available DASD 68 storage capacity generally impacts the degree to which a subscribing customer can effectuate VCR-type control over the presentation of a selected multimedia program. As illustrated... ..video segments 48 comprising the two-hour movie is transmitted only once from the multimedia server 30 to the subscriber's set-top control system 62. Moving outside of the presentation... ..compressed video segments 48. Such incidents of re-transmission preferably result in additional costs being charged to the subscriber's account.

With further reference to Fig. 11, the set-top controller 64 of the set-top control system 62 preferably communicates with a remote multimedia server 30 over a communication channel 44, and coordinates the operation of the set-top control system 62. Media-on-demand data is generally transmitted from the multimedia server 30 to the set-top control system 62 over the communication channel 44 at a... ..coordinate the reception, storage, and decoding of compressed video segments 48 received from the multimedia server 30, and the presentation of the decoded video segments 48 on a subscribing customer's television 76. The set-top controller 64 preferably communicates control signals to the multimedia server 30 over a server control line or channel 78 of the communication channel 44 to initiate transmission of a... ..regulate the rate at which the compressed video signal stream is received from the multimedia server 30 over the data channel 75 to avoid an input buffer 66 overflow condition.

During... ..a control signal is preferably issued by the set-top controller 64 to the multimedia server 30 over the server control line 78 to request temporary halting of source video signal stream transmission, thus causing... ..remain stationary. The set-top controller 64 preferably issues a resume control command over the server control line 78 when requesting the multimedia server 30 to resume transmission of the source video signal stream. By way of further example, a subscribing customer may view portions of the multimedia program outside of the presentation control window 90 by... ..alerts the subscriber that satisfying the request will require additional video data from the multimedia server 30 and result in an associated charge to the subscriber's account. A subscriber may initiate transmission of the additional video data... ..signals to the input buffer 66, DASD 68, output buffer 72, decoder 74, and multimedia server 30 to regulate timing and data transmission within the set-top control system 62 respectively... ..temporarily buffering video segments being transferred into and out of the DASD 68 to enhance synchronization, and to buffer information packets and other data unrelated to the video segment 48 data prior to being... ..number is preferably used as an identification address when routing video data from the multimedia server 30 to the set-top control system 68 of the subscribing customer who placed the pay-per-view order. As discussed previously hereinabove, an ATM information cell... ..be situated at an outer diameter disk location, an inner diameter disk location, or an intermediate diameter disk location. It is noted that a nominal disk 108 rotation rate should be... ..the actuator performs a seek to locate a new track, it must generally decelerate and settle to a position in which it is following the centerline of the data track. Generally, a longer period of time is required for the actuator to settle at the end of a seek operation for narrower track widths, thereby increasing the overall... ..operations are performed, and, as a result, the time required for the actuator 112 to settle is no longer a significant factor that might otherwise limit the degree to which the... ..Many read errors are often imperceivable to the viewing or listening observer. Moreover, various signal processing and smoothing techniques may be employed to enhance the audio and video presentation upon the... ..Kbpi (Kilobits per inch). It is noted that more data can be stored per linear unit of track length in a spiral data track in comparison to conventional concentric tracks due... ..in the previously identified related U.S. Patent Application Serial No. 08/288,525 entitled "Apparatus and Method for Providing Multimedia Data."

MULTIMEDIA DASD DATA STORAGE ARCHITECTURE

Local customized control over... ..accessing of non-sequentially ordered and sequentially ordered video segments 54 received from a multimedia server 30 preferably of the type previously described. For purposes of clarity and

simplicity of explanation... only, and do not represent limitations as to the scope of the disclosed method and apparatus.

Referring now to Figs. 15-19, it is assumed, for purposes of explanation, that the... window 90 for effectuating full VCR-type presentation control functions is twenty seconds, and the customer selected movie is two-hours in duration. It is noted that a typical presentation control... control system 62 is configured to store two discrete video segments 48. Accordingly, the multimedia server 30 transmits video segment packets containing no more than two video segments 48 to the... cost configuration. Such a low-cost configuration typically requires frequent packet transmissions from the multimedia server 30, thereby increasing the service costs associated with receiving multimedia programming from the multimedia server 30.

Further, it is assumed that the MPEG-1 compression standard is employed to obtain... 110 and 111 is respectively shown by the direction arrows provided in Fig. 19. This process of sweeping over one surface of the disk 108, performing a head switch operation, and... control window 90, thus requiring transmission of additional video segment 48 information from the multimedia

server 30.

PRESENTATION CONTROL WINDOW ARCHITECTURE

Still referring to Figs. 18 and 19 in detail, it... message annunciating the reception of an incoming communication from a source other than the multimedia server 30, for example, may be received by the input buffer 66 and transferred to the... video segment

T_0 = Decompressed full-motion program time in seconds per video segment

P = Maximum server packet size in number of video segments based on subscriber's input buffer capacity in... Input buffer of set-top control system is preferably configured to store at least two server packets (P) to allow server flexibility when asynchronously transmitting video segment packets (i.e., $IBS > 2 \times P \times S_0$... video segment 48. Assuming that the maximum size of each packet transmitted by the multimedia server 30 is two segments ($P = 2$), the set-top control system's 62 input buffer... having an average size of 0.167 MB, or approximately 5 MB. For increased multimedia server 30 asynchronous transmission flexibility, the input buffer should have a storage capacity of approximately 10... a multimedia communication system. The formatting of the multimedia information received from a remote multimedia server 30 may be varied in accordance with the operational characteristics, specifications, and functions of a... multimedia program. In one illustrative embodiment, the video segment 48 information transmitted by the multimedia server 30 to the set-top control system 62 in discrete packets shown in Fig. 10... video segments (S_0), for example, multimedia program information can be efficiently transmitted from the multimedia server 30 in a format specifically tailored to the system configuration and control functionality of a subscribing customer's unique set-top control system 62.

ASYNCHRONOUS FORMATTING METHODOLOGY

Turning now to Figs. 21... depicted in chart form in Figs. 20. A subscriber preferably communicates with a remote multimedia server 30 through a novel set-top control system 62 preferably of a type discussed in... pay-per-view basis, at step 300. In this example, it is assumed that the customer is interested in selecting among various video programs, such as feature-length movies. At step... art, or a point-and-click interface similar to that commonly used when communicating with computer systems, for example.

At step 304, the subscriber preferably specifies the duration or capacity of... as a fifty minute lecture previously recorded at a local university, for example, a subscribing customer may wish to specify a fifty minute duration for the presentation control window 90 so... system 62, among others.

After selecting one or more desired multimedia programs from the multimedia server menu, the set-top controller 64, at step 308, preferably performs various internal computations to determine the nominal DASD 68 storage capacity needed to support the customer-specified presentation control window 90. The predetermined time duration (PTD) specified by the subscriber is... associated with the presentation control window 90 and the configuration and functionality of a subscribing customer's set-top control system 62 are transmitted to the multimedia server 30. The multimedia server 30 preferably includes a server controller 34 that, at step 316, reads the configuration parameters received from a subscriber's... can accommodate the subscriber-specified presentation control window 90 may instead be performed by the server controller 34 based upon the received set-top control system 62 configuration parameters.

As previously discussed, a selected multimedia program may be stored in the multimedia server 30 in either an analog format or a digital format. A selected multimedia program stored... a local, national, or international network broadcast channel 45, is typically received by the multimedia server 30 in an analog format, and may also be digitized at step 318. The digitized... segmentizing operations of steps 318 and 320 are typically not applicable to multimedia programs previously processed and stored in the multimedia server 30 in a digital format. These steps are preferably performed only once when initially storing a multimedia program on a digital storage device 35 within the multimedia server 30.

Turning now to Fig. 22 the sequential program segments comprising the selected multimedia program are preferably arranged in a customized order at step 330. In one embodiment, the multimedia server 30 includes a video parser 38 that preferably transforms the sequential program segments into a... 62. In response to a subscriber's configuration parameters, the controller 34 of the multimedia server 30 preferably determines the number of segment blocks (M) per disk surface (D), from which... derived the Block Indexing Coefficient, $BI = \text{modulo } D \times M$, at step 332. Further, the server controller 34, at step 334, preferably determines the length (L) of each of the segment... of rows of each customized matrix comprising a segment block (M).

At step 336, the server controller 34 preferably computes the size of the program segment packets, which is typically dependent on the size (IBS) of the input buffer 66 of a subscribing customer's set-top control system 62. The input buffer 66 of a set-top control... is transmitted to a particular set-top control system 62, is preferably computed by the server controller 34 in a manner previously discussed hereinabove. At step 340, the program segments 48...68 formatting methodology. The set-top control system 62 formatting parameters transmitted by the multimedia server 30 are received by the subscriber's set-top control system 62 at step 350... parameters provide information preferably used by the set-top controller 64 to properly buffer and process the customized program segment packets received over the communication channel 44. For example, the segment... in turn, preferably coordinates writing and reading of the program segments received from the multimedia server 30 to and from a corresponding number (M) of storage blocks of a predetermined length... 108 surface.

At step 351, a customized program segment 48 series transmitted from the multimedia server 30 in packets over the communication channel 44 are received by the subscriber's set-top control system 62. Generally, one packet is received during each multimedia server 30 transmission window, although multiple packets may be transmitted during this time if the input... discussed, concerns the concurrent buffering and displaying of program segments 48 received from the multimedia server 30 to facilitate virtually instantaneous on-demand viewing of a selected multimedia program. As previously... 358. It is noted that, at steps 356 and 358, it may be advantageous for synchronization purposes to transfer a number of non-sequential program segments 48 received from the communication... have been transferred to the output buffer 76 in sequential order.

As mentioned previously, the process of concurrently transferring non-sequential program segments 48 to both the DASD 68, at step... 74 thus provides for virtually instantaneous presentation of a selected multimedia program on a subscribing customer's television 76. The program segments 48 contained in subsequently received Packet-3 and Packet... and, if an overflow condition is imminent, preferably transmits a control signal to the multimedia server 30 to request temporary halting of the transmission of program segment packets at step 362... periods of normal

program viewing since the transmission and reception of program segment packets is synchronized by transferring packets during prescribed transmission windows. Various presentation control window 90 function modes, such...
 ...of a halt control signal from the set-top control system 62 to the multimedia server 30. The remaining program segments 48 contained in Packet-1 and stored in the input... ...segment A12 has been read from Location-18 during RUN-3 and displayed on the customer's television 74, and that the operations described in Figs. 24-26 are associated with... ...is read from Location-4, and decoded and then displayed at step 410 on the customer's television 76 in sequence with respect to the previously read and displayed program segment... ...is read from Location-14, and decoded and then displayed at step 430 on the customer's television 76 in sequence with respect to the previously read program segment A14. The...

6/K/24 (Item 24 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Patent Assignee:

- TANDEM COMPUTERS INCORPORATED... ;
 ; ;

Country	Number	Kind	Date
---------	--------	------	------

Abstract ...A2

Type	Pub. Date	Kind	Text
Available Text	Language	Update	Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: ...and 08/485,053, filed concurrently herewith.

The present invention is directed generally to data processing systems, and more particularly to a multiple processing system and a reliable system area network that provides connectivity for interprocessor and input/output... ...and communications systems to general purpose high availability commercial systems. The evolution of fault tolerant computers has been well documented (see D. P. Siewiorek, R. S. Swarz, "The Theory and Practice... ...and the Jet Propulsion laboratory began to apply fault tolerance to the development of guidance computers for aerospace applications. The 1960's also saw the development of the first AT&T electronic switching systems.

The first commercial fault tolerant machines were introduced by Tandem Computers in the 1970's for use in on-line transaction processing applications (J. Bartlett, "A NonStop Kernal," in proc. Eighth Symposium on Operating System Principles, pp... ...systems were introduced in the 1980's (O. Serlin, "Fault- Tolerant Systems in Commercial Applications," Computer, pp. 19-30, August 1984). Current commercial fault tolerant systems include distributed memory multi-processors, shared-memory transaction based systems, "pair-and-spare" hardware fault tolerant systems (see R. Freiburghouse, "Making Processing Fail-safe," Mini-micro Systems, pp. 255-264, May 1982; U.S. Patent No. 4... ...system.), and triple-modular-redundant systems such as the "Integrity" computing system

manufactured by Tandem Computers Incorporated of Cupertino, California, assignee of this application and the invention disclosed herein.

Most applications of commercial fault tolerant computers fall into the category of on-line transaction processing. Financial institutions require high availability for electronic funds transfer, control of automatic teller machines, and... ..telecommunications systems.

Vendors of fault tolerant machines attempt to achieve both increased system availability, continuous processing, and correctness of data even in the presence of faults. Depending upon the particular system architecture, application software ("processes") running on the system either continue to run despite failures, or the processes are automatically restarted from a recent checkpoint when a fault is encountered. Some fault tolerant systems are provided with sufficient component redundancy to be able reconfigure around failed components, but processes running in the failed modules are lost. Vendors of commercial fault tolerant systems have extended fault tolerance beyond the processors and disks. To make large improvements in reliability, all sources of failure must be addressed... ..power supplies, fans and inter-module connections.

The "NonStop," and "Integrity" architectures manufactured by Tandem Computers Incorporated, (both respectively illustrated broadly in U.S. Patent No. 4,228,496 and U... ..assigned to the assignee of this application; NonStop and Integrity are registered trademarks of Tandem Computers Incorporated) represent two current approaches to commercial fault tolerant computing. The NonStop system, as generally... ..above-identified U.S. Patent No. 4,278,496, employs an architecture that uses multiple processor systems designed to continue operation despite the failure of any single hardware component. In normal operation, each processor system uses its major components independently and concurrently, rather than as "hot backups". The NonStop system architecture may consist of up to 16 processor systems interconnected by a bus for interprocessor communication. Each processor system has its own memory which contains a copy of a message-based operating system. Each processor system controls one or more input/output (I/O) busses. Dual-porting of I/O controllers and devices provides multiple paths to each device. External storage (to the processor system), such as disk storage, may be mirrored to maintain redundant permanent data storage.

This... ..hardware, while fault recovery is the responsibility of the software.

Also, in the Nonstop multi-processor architecture, application software ("process") may run on the system under the operating system as "process-pairs," including a primary process and a backup process. The primary process runs on one of the multiple processors while the backup process runs on a different processor. The backup process is usually dormant, but periodically updates its state in response to checkpoint messages from the primary process. The content of a checkpoint message can take the form of complete state update, or currently most application code runs under transaction processing software which provides recovery through a combination of checkpoints and transaction two-phase commit protocols.

Interprocessor message traffic in the Tandem Nonstop architecture includes each processor periodically broadcasting an "I'm Alive" message for receipt by all the processors of the system, including itself, informing the other processors that the broadcasting processor is still functioning. When a processor fails, that failure will be announced and identified by the absence of the failed processor's periodic "I'm Alive" message. In response, the operating system will direct the appropriate backup processes to begin primary execution from the last checkpoint. New backup processes may be started in another processor, or the process may be run with no backup until the hardware has been repaired. U.S. Patent... ..example of this technique.

Each I/O controller is managed by one of the two processors to which it is attached. Management of the controller is periodically switched between the processors. If the managing processor fails, ownership of the controller is automatically switched to the other processor. If the controller fails, access to the data is maintained through another controller.

In addition to providing hardware fault tolerance, the processor pairs of the above-described architecture provide some measure of software fault tolerance. When a processor fails due to a software error, the backup processor frequently is able to successfully continue processing without encountering the same error. The software environment in the backup processor typically has different queue lengths, table sizes, and process mixes. Since most of the software bugs escaping the software quality assurance tests involve infrequent data dependent boundary conditions, the backup processes often succeed.

In contrast to the above-described architecture, the Integrity system illustrates another approach... fault recovery is the logical choice since few modifications to the software are required. The processors and local memories are configured using triple-modular-redundancy (TMR). All processors run the same code stream, but clocking of each module is independent to provide tolerance... three streams is asynchronous, and may drift several clock periods apart. The streams are re-synchronized periodically and during access of global memory. Voters on the TMR Controller boards detect and mask failures in a processor module. Memory is partitioned between the local memory on the triplicated processor boards and the global memory on the duplicated TMRC boards. The duplicated portions of the... techniques to detect failures. Each global memory is dual ported and is interfaced to the processors as well to the I/O Processors (IOPs). Standard VME peripheral controllers are interfaced to a pair of busses through a Bus... the BIMs to switch control of all controllers to the remaining IOP. Mirrored disk storage units may be attached to two different VME controllers.

In the Integrity system all hardware failures... reintegrated on-line.

The preceding examples illustrate present approaches to incorporating fault tolerance into data processing systems. Approaches involving software recovery require less redundant hardware, and offer the potential for some... have been developed on other systems.

Thus, the systems described above provide fault tolerant data processing either by hardware (e.g., fail-functional, employing redundancy) or by software techniques (fail-fast... hardware). However, none of the systems described are believed capable of providing fault tolerant data processing, using both hardware (fail-functional) and software (fail-fast) approaches, by a single data processing system.

Computing systems, such as those described above, are often used for electronic commerce: electronic data interchange (EDI) and global messaging. Today's demands upon such electronic commerce, however, is demanding more and more throughput capacity as the number of users increases and... networks such as local area networks (LANS), and the like.

A key requirement for a server architecture is the ability to move massive quantities of data. The server should have high bandwidth that is scalable, so that added throughput capacity can be added... response time, latency affects service levels and employee productivity.

The present invention provides a multiple-processor system that combines both of the two above-described approaches to fault tolerant architecture, hardware redundancy and software recovery techniques, in a single system.

Broadly, the present invention includes a processing system composed of multiple sub-processing systems. Each sub-processing system has, as the main processing element, a central processing unit (CPU) that in turn comprises a pair of processors operating in lock-step, synchronized fashion... execute each instruction of an instruction stream at the same time. Each of the sub-processing systems further include an input/output (I/O) system area network system that provides redundant communication paths between various components of the larger processing

system, including a CPU and assorted peripheral devices (e.g., mass storage

units, printers, and the like) of a sub-processing system, as well as between the sub-processors that may make up the larger overall processing system. Communication between any component of the processing system (e.g., a

CPU and another CPU, or a CPU and any peripheral device, regardless of which sub-processing system it may belong to) is implemented by forming and transmitting packetized messages that are... responsible for choosing the proper or available communication paths from a transmitting component of the processing system to a destination component based upon information contained in the message packet. Thus, the... peripherals, but permits it to also be used for interprocessor communications.

As indicated above, the processing system of the present invention is structured to provide fault-tolerant operation through both "fail... at a variety of points in the various data paths between the (lock-step operated) processor elements of the CPU and its associated memory. In particular, the processing system of the present invention conducts error-checking at an interface, and in a manner... little impact on performance. Prior art systems typically implement error-checking by running pairs of processors, and checking (comparing) the data and instruction flow between the processors and a cache memory. This technique of error-checking tended to add delay to the... error-checking precluded use of off-the-shelf parts that may be available (i.e., processor/cache memory combinations on a single semiconductor chip or module). The present invention performs error-checking of the processors at points that operate at slower rates, such as the main memory and I/O interfaces which operate at slower speeds than the processor-cache interface. In addition, the error-checking is performed at locations that allow detection of errors that may occur in the processors, their cache memory, and the I/O and memory interfaces. This allows simpler designs for... other data integrity checks.

Error-checking of the communication flow between the components of the processing system is achieved by adding a cyclic-redundancy-check (CRC) to the message packets that... "Good" (TPG) or "This Packet Bad" (TPB) - is appended to every packet. A maintenance diagnostic processor can use this information to isolate a link or router element that introduces an error... of topologies, so that alternate paths can be provided between any two elements of a processing system (e.g., between a CPU and an I/O device), for communication in the... so (e.g., by creating a "deadlock" condition, discussed further below).

The CPUs of a processing system are capable of operating in one of two basic modes: a "simplex mode" in... independently of the other, or a "duplex" mode in which pairs of CPUs operate in synchronized, lock-step fashion. Simplex mode operation provides the capability of recovering from faults that are... U.S. Pat. No. 4,228,496 which teaches a multiprocessing system in which each processor has the capability of checking on the operability of its sibling processors, and of taking over the processing of a processor found or believed to have failed). When operating in duplex mode, the paired CPUs both... fault tolerant platform for less robust operating systems (e.g., the UNIX operating system). The processing system of the present invention, with the paired, lock-step CPUs, is structured so that... masked (i.e., operating despite the existence of a fault), primarily through hardware.

When the processing system is operating in duplex mode, each CPU pair uses the I/O system to access any peripheral of the processing system, regardless of which (of the two, or more) sub-processor system the peripheral may be ostensibly a member of. Also, in duplex mode, message packets... message for the CPU pair (from either a peripheral device such as a mass storage unit or from a processing unit), will replicate the message and deliver it to both CPUs of the pair using synchronization methods that ensure that the CPUs remain synchronized. In effect, the duplex CPU pair, as viewed from the I/O system and other... as a single CPU. Thus, the I/O system, which includes elements from all sub-processing systems, is made to be seen by the duplex CPU pair as one homogeneous system... a multiprocessor system in which the CPU of any one is actually a pair of synchronized, lock-step CPUs.

Yet another important aspect of the present invention is that interrupts issuing... interrupts via the message packet system ensures that they will arrive at duplexed CPUs in synchronized fashion, in the same manner as I/O message packets. Interrupt message packets will contain... the system. In addition, using the same messaging system to communicate data between I/O units and the CPUs and to communicate interrupts to the CPUs preserves the ordering of I... the implementation of a technique of validating access to the memory of any CPU. The processing system, as structured according to the present invention, permits the memory of any CPU to a CPU and any other component of the processor system. Thereby, the individual processor units of the CPU are removed from the more

mundane tasks of getting information from memory and out onto the TNet network, or accepting information from the network. The processor unit of the CPU merely sets up data structures in memory containing the data to be...
...is required, where in memory the response is to be placed when received. When the processor unit completes the task of creating the data structure, the block transfer engine is notified to...
...response is received, it is routed to the expected memory location identified, and notifies the processor unit that the response was received.

Further aspects and features of the present invention will become...
...invention, which should be taken in conjunction with the accompanying drawings.

Fig. 1A illustrates a processing system constructed in accordance with the teachings of the present invention, and Figs. 1B and 1C illustrate two alternate configurations of the processing system of Fig. 1A, employing clusters or arrangements of the processing system of Fig. 1A;

Fig. 2 illustrates, in simplified block diagram form, the central processing unit (CPU) that forms a part of each sub-processor system of Figs. 1A - 1C;

Figs. 3A - 3D and 4A - 4C illustrate the construction of...
...the area network I/O system shown in Fig. 2;

Fig. 5 illustrates the interface unit that forms a part of the CPUs of Fig. 2 to interface the processor and memory with the I/O area network system;

Fig. 6 is a block diagram, illustrating a portion of packet receiver of the interface unit of Fig. 5;

Fig. 7A diagrammatically illustrates the clock synchronization FIFO (CS FIFO) used by the packet receiver section packet receiver shown in Fig. 6;

Fig. 7B is an block diagram of a construction of the clock synchronization FIFO structure shown in Fig. 7A;

Fig. 8 illustrates the cross-connections for error-checking outbound transmissions from the two interface units of a CPU;

Fig. 9 illustrates an encoded (8B to 9B) data/command symbol;

Fig. 10 illustrates the method and structure used by the interface unit of Fig. 5 to cross-check for errors data being transferred to the memory controllers...
...of a CPU of Fig. 2 to other (external to the CPU) components of the processing system;

Fig. 12 is a block diagram that diagrammatically illustrates the formation of an address...
...14A illustrates the logic for posting interrupt requests to queues in memory and to the processor units of the CPU of Fig. 2;

Fig. 14B illustrates the process used to form a memory address for a queue entry;

Fig. 15 is a block...
...data output constructs formed in the memory of the CPU of Fig. 2 by a processor unit, and containing data to be sent via the area I/O networks shown in Figs. 1A - 1C, and also illustrating the block transfer engine (BTE) unit of the interface unit of Fig. 5 that operates to access the data output constructs for transmission to the...
...pair of memory controllers between memory of a CPU of Fig. 2 and its interface unit for accessing from memory 72 bits of data, including two simultaneously-accessed 32-bit words...
...other for error-checking;

Fig. 19A is a simplified block diagram illustration of the router unit used in the area input/output networks of the processing systems shown in Figs. 1A - 1C;

Fig. 19B illustrates comparison on two port inputs of the router unit of Fig. 19A;

Fig. 20A is a block diagram the construction of one of the six input ports of the router unit shown in Fig. 19A;

Fig. 20B is a block diagram of the synchronization logic used to validate command/data symbols received at an input port of the router unit of Fig. 19A;

Fig. 21A is a block diagram illustration of the target port selection... ..is a block diagram illustration of one of the six output ports of the router unit shown in Fig. 19A;

Fig. 23 is an illustration of the method used to transmit identical information to a duplexed pair CPUs of Fig. 2 in synchronized fashion when the processing system is operating in lock-step (duplex) mode, using a pair the FIFOs of Fig... ..is a simplified block diagram illustrating the clock generation system of each of the sub-processing systems of Figs. 1A - 1C for developing the plurality of clock signals used to operate the various elements of that sub-processing system;

Fig. 25 illustrates the topology used to interconnect the clock generation systems of paired sub-processing systems for synchronizing the various clock signals of the pair of sub-processing systems to one another;

Fig. 26A and 26B illustrates a FIFO constant rate clock control logic used to control the clock synchronization FIFO of Figs. 8 or 20 in the situation when the two clocks used to... ..structure of the on-line access port (OLAP) used to provide access to the maintenance

processor (MP) to the various elements of the system of Fig. 1A (or those of Figs... ..the soft-flag logic used to handle asymmetric variables between the CPUs of paired sub-processing systems operating in duplex mode;

Fig. 31A shows a flow diagram, and Fig. 31B illustrates a portion of SYNC CLK, both of which are used to reset and synchronize the clock synchronization FIFOs of the CPUs and routers of the processing system of Fig. 1A that receive information from each other;

Fig. 32 is a flow... ..33A - 33D generally illustrate the procedure used to bring an one of the CPUs of processing system shown in Fig. 1A into lock-step, duplex mode operation with the other of the CPUs without measurably halting operation of the processing system; and

Fig. 34 illustrates a reduced cost architecture incorporating teachings of the invention; and... ..to the figures and, for the moment, principally Fig. 1A, there is illustrated a data processing system, designated with the reference 10, constructed according to the various teachings of the present invention. As Fig. 1A shows, the data processing system 10 comprises two sub-processor systems 10A and 10B each of which are substantially the same in structure and function... ..of any one of the subprocessor systems 10 will apply equally to any other sub-processor system 10.

Continuing with Fig. 1A therefore, each of the sub-processor systems 10A, 10B is illustrated as including a central processing unit (CPU) 12, a router 14, and a plurality of input/output (I/O) packet interfaces... ..one of the I/O packet interfaces 16 will also have coupled thereto a maintenance processor (MP) 18.

The MP 18 of each sub-processor system 10A, 10B connects to each of the elements of that sub-processor system via an IEEE 1149.1 test bus 17 (shown in phantom in Fig. 1A... ..accompanying clock signal. As Fig. 1A further illustrates, TNet Links L also interconnect the sub-processor systems 10A and 10B to one another, providing each sub-processor system 10 with access to the I/O devices of the other as well as inter-CPU communication. As will be seen, any CPU 12 of the processing system 10 can be given access to the memory of any other CPU 12, although... ..the memory of a CPU 12 by a wayward peripheral device 17.

Preferably, the sub-processor systems 10A/10B are paired as illustrated in Fig. 1A (and Figs 1B and 1C, discussed below), and each sub-processor system 10A/10B pair (i.e., comprising a CPU 12, at least one router 14... ..12A) connects, by a TNet Link L to a router (14A) of the corresponding sub-processor system (e.g., 10A). Conversely, the Y port connects the CPU (12A) to the router (14B) of the companion sub-processor system (10B). This latter connection not only provides a communication path for access by a CPU (12A) to the I/O devices of the other sub-processor system (10B), but also to the CPU (12B) of that system for inter-CPU communication.

Information is communicated between any element of the processing system 10 and any other element (e.g., CPU 12A of sub-processor system 10A) of the system and any other element of the system (e.g., an I/O device associated with an I/O packet interface 16B of sub-processor system 10B) via message "packets." Each message packet is made up of a number of... this reason, a unique method of receiving the symbols at the receiver, using a clock synchronization first-in-first-out (CS FIFO) storage structure (described more fully below), has been developed... operation means just that: the frequencies of the clock signals of the transmitter and receiver units are locked, although not necessarily in phase. Frequency locked clock signals are used to transmit symbols between the routers 14A, 14B and the CPUs 12 of paired sub-processor systems (e.g., sub-processor systems 10A, 10B, Fig. 1A). Since the clocks of the transmitting and receiving element are not phase related, a clock synchronization FIFO is again used -- albeit operating in a slightly different mode from that used for... difference, as will be seen, is due to the fact that pairs of the sub-processor systems 10 can be operated in a synchronized, lock-step mode, called duplex mode, in which each CPU 12 operates to execute the... 1A illustrates another feature of the invention: a cross-link connection between the two sub-processor systems 10A, 10B through the use of additional routers 14 (identified in Fig. 1A as... RY(sub(1)), and RY(sub(2)) form a cross-link connection between the sub-processors 10A, 10B (or, as shown, "sides" X and Y, respectively) to couple them to I... the routers RX(sub(2)) and RY(sub(2)) provide the I/O packet interface units 16x and 16y with a dual ported interface. Of course, it will now be evident... lend themselves to being used in a manner that can extend the configuration of the processing system 10 to include additional sub-processor systems such as illustrated in Figs. 1B and 1C. In Fig. 1B, for example, one... of each of the routers 14A and 14B is used to connect the corresponding sub-processor systems 10A and 10B to additional sub-processor systems 10A' and 10B' forming thereby a larger processing system comprising clusters of the basic processing system 10 of Fig. 1.

Similarly, in Fig. 1C the above concept is extended to form an eight sub-processor system cluster, comprising sub-processor systems pairs 10A/10B, 10A'/10B', 10A''/10B'', and 10A'''/10B'''. In turn, each of the sub-processor systems (e.g., sub-processor system 10A) will have essentially the same basic minimum configuration of a CPU 12, a... by a I/O packet interface 16, except that, as Fig. 1C shows, the sub-processor systems 10A and 10B include additional routers 14C and 14D, respectively, in order to extend the cluster beyond sub-processor systems 10A'/10B' to the sub-processor systems 10A''/10B'' and 10A'''/10B''''. As Fig. 1C further illustrates, unused ports 4 and... the routers 14 when configuring the topology of the system 10, any CPU 12 of processing system 10 of Fig. 1C can access any other "end unit" (e.g., a CPU or I/O device) of any of the other sub-processor systems. Two paths are available from any CPU 12 to the last router 14 connecting to the I/O packet interface 16. For example, the CPU 12B of the sub-processor system 10B' can access the I/O 16''' of sub-processor system 10A''' via router 14B (of sub-processor system 10B'), router 14D, and router 14B (of sub-system 10B''') and, via link LA... 10A'''), OR via router 14A (of sub-system 10A'), router 14C, and router 14A (sub-processor system 10A'''). Similarly, CPU 12A of sub-processor system 10A'' may access (via two paths) memory contained in the CPU 12B of sub-processor 10B to read or write data. (Memory accesses by one CPU 12 of another component of the processing system requires, as will be seen, the components seeking access to have authorization to do... prevents corruption of memory data of a CPU by erroneous access.)

The topology of the processing system shown in Fig. 1B is achieved by using port 1 of the routers 14A, 14B, and auxiliary TNet links LA, to connect to the routers 14A', 14B' of sub-processor systems 10A', 10B'. The topology thereby obtained establishes redundant communication paths between any CPU 12 (12A, 12B, 12A', 12B') and any I/O packet interface 16 of the processing system 10 shown in Fig. 1B. For example, the CPU 12A' of the sub-processor system 10A' may access the I/O 16A of sub-processor system 10A by a first path formed by the router 14A' (in port 4, out... shown in Fig. 1B. By interconnecting one port of each router 14 of each sub-processor pair, and using additional auxiliary TNet links LA (illustrated in Fig. 1C with the dotted line connections) between the ports 1 of the routers 14 (14A'' and 14B'') of sub-processor systems 10A'', 10B'' and 10A''', 10B''', two separate, independent data paths can be found between any CPU 12 and any I/O packet interface 16. In this fashion, any end unit (i.e., a CPU 12 or an I/O packet interface 16) will have at least two paths to any other end unit.

Providing alternate paths of access between any two end units (e.g., between a CPU 12 and any other CPU 12, or between any CPU... any two of the remaining fault domains. Here, a fault domain could be a sub-processor system (e.g., 10A). Thus, if the sub-processor system 10A were brought down because of a failure the electrical power being supplied, without... TNet link LA between the routers 14A' and 14B'', the CPU 12B of the sub-processor system 10B would have lost access to the I/O packet interface 16'' (via router... with the loss of the router 14A (and router 14C) by loss of the sub-processor system 10A, communications between the CPU 12B is still possible via the route of router... equally to CPU 12B. As Fig. 2 shows, the CPU 12A includes a pair of processor units 20a, 20b that are configured for synchronized, lock-step operation in that both processor units 20a, 20b receive and execute identical instructions, and issue identical data and command outputs, at substantially the same moments in time. Each of the processor units 20a and 20b is connected, by a bus 21 (21a, 21b) to a corresponding cache memory 22. The particular type of processor units used could contain sufficient internal cache memory so that the cache memory 22 would not... 22 could be used to supplement any cache memory that may be internal to the processor units 20. In any event, if the cache memory 22 is used, the bus 21 is... 22 address bits, 3 bits of parity covering the address, and 7 control bits.

The processors 20a, 20b are also respectively coupled, via a separate 64-bit address/data bus 23 to X and Y interface units 24a, 24b. If desired, the address/data communicated on each bus 23a, 23b could also be protected by parity, although this will increase the width of the bus. (Preferably, the processors 20 are constructed to include RISC R4000 type microprocessors, such as are available from the MIPS Division of Silicon Graphics, Inc. of Santa Clara, California.)

The X and Y interface units 24a, 24b operate to communicate data and command signals between the processor units 20a, 20b and a memory system of the CPU 12A, comprising a memory controller (MC... MC halves 26a and 26b) and a dynamic random access memory array 28. The interface units 24 interconnect to each other and to the Mcs 26a, 26b by a 72-bit... accompanied by 8 bits of ECC) are written to the memory 28 by the interface units 24, one interface unit 24 will drive only one word (e.g., the 32 most significant portion) of the doubleword being written while the other interface unit 24 writes the other word of the double word (e.g., the least significant 32-bit portion of the doubleword). In addition, on each write operation the interface units 24a, 24b perform a cross-check operation on the data not written by that interface unit 24 with the data written by the other to check for errors; on read operations... accessed corresponds to the address of the location from which the doubleword was stored.

Interface units 24a, 24b of the CPU 12A form the circuitry to respectively service the X and Y (I/O) ports of the CPU 12A. Thus, the X interface unit 24a connects by the bi-directional TNet Link Lx to a port of the router 14A of the processor system 10A (Fig. 1A) while the Y interface unit 24b similarly connects to the router 14B of the processor system 10B by TNet Link Ly. The X interface unit 24a handles all I/O traffic between the router 14A and the CPU 12A of the sub-processor system 10A. Likewise, the Y interface unit 24b is responsible for all I/O traffic between the CPU 12A and the router 14B of companion sub-processor system 10B.

The TNet Link Lx connecting the X interface unit 24a to the router 14A (Fig. 1) comprises, as above indicated, two 10-bit buses... sub(x)) carries data incoming from the router 14A. In similar fashion, the Y interface unit 24b is connected to the router 14B (of the sub-processor system 10B) by two 10-bit busses: 30(sub(y)) (for outgoing transmissions) and 32... y)) (for incoming transmissions), together forming the TNet Link Ly.

The X and Y interface units 24a, 24b are synchronously operated in lock-step, performing substantially the same operations at substantially the same times. Thus, although only the X interface unit 24a actually transmits data onto the bus 30(sub(x)), the same output data is being produced by the Y interface unit 24b, and used for error-checking. The Y interface unit 24b output data is coupled to the X interface unit 24a by a cross-link 34(sub(y)) where it is received by the X interface unit 24a and compared against the same output data produced by the X interface unit. In this way the outgoing data made available at the X port of the CPU... the port of the CPU 12A is checked. The

output data from the Y interface unit 24b is coupled to the Y port by a 10-bit bus 30(sub(y)), and also to the X interface unit 24a by the 9-bit cross-link 34(sub(y)) where is checked with that produced by the X interface unit.

As mentioned, the two interface units 24a, 24b operate in synchronous, lock-step with one another, each performing substantially the same... ..X and/or Y ports of the CPU 12A must be received by both interface units 24a, 24b to maintain the two interface units in this lock-step mode. Thus, data received by one interface unit 24a, 24b is passed to the other, as indicated by the dotted lines and 9... ..sub(x)) (communicating incoming data being received at the X port by the X interface unit 24a to the Y interface unit 24b) and 36(sub(y)) (communicating data received at the Y port by the Y interface unit 24b to the X interface unit 24a).

Certain more robust operating systems are structured with a fault-tolerant capability in the... ..example, U.S. Patent No. 4,817,091 teaches a multiprocessor system in which each processor periodically messages each of the processors of the system (including itself), under software control, to thereby provide an indication of continuing operation. Each of the processors, in addition to performing its normal tasks, operates as a backup processor to another of the processors. In the event one of the backup processors fails to receive the messaged indication from a sibling processor, it will take over the operation of that sibling (now thought to be inoperative), in... ..platform for both types of software. Thus, when a robust operating system is available, the processing system 10 can be configured to operate in a "simplex" mode in which each of... ..left, in most instances, to software.

Alternatively, for less robust operating systems and software, the processing system 10 provides a hardware-based fault-tolerance by being configured to operate in a... ..g., CPUs 12A, 12B) are coupled together as shown in Fig. 1A, to operate in synchronized, lock-step fashion, executing the same instructions at the substantially the same moment in time... ..data and command symbols. In order to simplify the design of the CPU 12, the processors 20 are precluded from communicating directly with any outside entity (e.g., another CPU 12... ..O device via the I/O packet interface 16). Rather, as will be seen, the processor will construct a data structure in memory and turn over control to the interface units 24. Each interface unit 24 includes a block transfer engine (BTE; Fig. 5) configured to provide a form of... ..to the destination according to information contained in the message packet.

The design of the processing system 10 permits a memory 28 of a CPU to be read or written by... ..via the routers 14. Accordingly, before continuing with the description of the construction of the processing system 10, it would be of advantage to understand first the configuration of the data...information.

As indicated, the HADC message packet operates to communicate write data between the end units (e.g., CPU 12) of the processing system 10. Other message packets, however, may be differently constructed because of their function and... ..CRC. The HC message packet is used to acknowledge a request to write data.

Interface Unit:

The X and Y interface units 24 (i.e., 24a and 24b - Fig. 2) operate to perform three major functions within the CPU 12: to interface the processors 20 to the memory 28; to provide an I/O service that operates transparently to, but under the control of, the processors; and to validate requests for access to the memory 28 from outside sources.

Regarding first the interface function, the X and Y interface units 24a, 24b operate to respectively communicate processors 20a, 20b to the memory controllers (Mcs 26a, 26b) and memory 28 for writing and... ..fast checking of the data read/written. For example, write operations have the two interface units 24a, 24b cooperating to cross-check the data to be written to ensure its integrity (and at the same time, the interface units 24 will operate) to develop an error correcting code (ECC) that covers, as will be With respect to I/O access, the processors 20 are not provided with the ability to communicate directly with the input/output systems... ..must write data structures to the memory 28 and then pass control to the interface units 24 which perform a direct memory access (DMA) operation to retrieve those data structures, and... ..indicated in the data structure itself.)

The third function of the X and Y interface units 24, access validation to the memory 28, uses an address validation and translation (AVT) table maintained by the interface units. The AVT table contains an address for each system

component (e.g., an I/O... ..the incoming message packets are virtual addresses. These virtual addresses are translated by the interface unit to physical addresses recognizable by the memory control units 26 for accessing the memory 28.

Referring to Fig. 5, illustrated is a simplified block diagram of the X interface unit 24a of the CPU 12A. The companion Y interface unit 24b (as well as the interface units 24 of the CPU 12B, or any other CPU 12) is of substantially identical construction. Accordingly, it will be understood that a description of the interface unit 24a will apply equally to the other interface units 24 of the processing system 10.

As Fig. 5 illustrates, the X interface unit 24a includes a processor interface 60, a memory interface 70, interrupt logic 86, a block transfer engine (BTE) 88, access validation and translation logic 90, a packet transmitter 94, and a packet receiver 96.

Processor Interface:

The processor interface 60 handles the information flow (data and commands) between the processor 20a and the X interface unit 24a. A processor bus 23, including a 64 bit address and data bus (SysAD) 23a and a 9 bit command bus 23b, couples the processor 20a and the processor interface 60 to one another. While the SysAD bus 23a carries memory address and data... ..and qualifying commands carried at substantially the same time on the SysAD bus 23a. The processor interface 60 operates to interpret commands issued by the processor unit 20a in order to pass reads/writes to memory or control registers of the processor interface. In addition, the processor interface 60 contains temporary storage (not shown) for buffering addresses and data for access to... ..26). Data and command information read from memory is similarly buffered en route to the processor unit 20a, and made available when the processor unit is ready to accept it. Further, the processor interface 60 will operate to generate the necessary interrupt signalling for the X interface unit 24a.

The processor interface 60 is connected to a memory interface 70 and to configuration registers 74 by a bit-directional 64 bit processor address/data bus 76. The configuration registers 74 are a symbolic representation of the various control registers contained in other components of the X interface unit 24a, and will be discussed when those particular components are discussed. However, although not specifically... ..throughout other of the logic that is used to implement the X interface 24a, the

processor address/data bus 76 is likewise coupled to read or write to those registers.

Configuration registers 74 are read/write accessible to the processor 20a; they allow the X interface unit to be "personalized." For example, one register identifies the node address of the CPU 12A... ..with the CPU 12A; another, readable only, contains a fixed identification number of the interface unit 24, and still other registers define areas of memory that can be used by, for... ..logic 90, etc.) employing them are discussed.

The memory interface 70 couples the X interface unit 24a to the memory controllers 26 (and to the Y interface unit 24b; see fig. 2) by a bus 25 that includes two 36 bi-directional bit... ..25a, 25b. The memory interface operates to arbitrate between requests for memory access from the processor unit 20, the BTE 88, and the AVT logic 90. In addition to memory accesses from the processor unit 20a, the memory 28 may also be accessed by components of the processing system 10 to, for example, store data requested to be read by the processor unit 20a from an I/O unit 17, or memory 28 may also be accessed for I/O data structures previously set up in memory by the processor unit. Since these accesses are all asynchronous, they must be arbitrated, and the memory interface 70... ..command information accessed from the memory 28 is coupled from the memory interface to the processor interface 60 by a memory read bus 82, as well as to an interrupt logic... ..doubleword quantities. However, while the memory interfaces 70 of both the X and Y interface units 24a and 24b ...by the memory interface 70 are coupled to the memory interface by the companion interface unit 24 where they are compared with the same 32 bits for error.

Digressing for the... ..containing interrupt information are received, that information is conveyed to the interrupt logic 86 for processing and posting for action by the processor 20, along with any interrupts generated internal to the CPU 12A. Internally generated interrupts will... ..register 71 (internal to the interrupt logic 86), indicating the cause of the interrupt. The processor 20 can then read and act upon the interrupt. The interrupt logic is discussed more fully below.

The BTE 88 of the X interface unit 24a operates to perform direct memory accesses, and provides the mechanism that allows the processors 20 to access external resources. The BTE 88 can be set-up by the processors 20 to generate I/O requests, transparent to the processors 20 and notify the processors when the requests are complete. The BTE logic 88 is discussed further below.

Requests for... ..8 byte wide format necessary for storing in the memory 28.

Outgoing message packets containing processor originated transaction requests (e.g., a read request asking for a block data from an I/O unit) are monitored by the request transaction logic (RTL) 100. The RTL 100 provides a time... ..will generate an interrupt (handled and reported by the interrupt logic 86) to inform the processor 20 that the request was not honored. In addition, the RTL 100 will validate responses... ..28 (by the DMA operation of the BTE 86) at a location known to the processor 20 so that it can locate the response.

Each of the CPUs 12 are checked... ..discussed. One such check is an on-going monitor of the operation of the interface units 24a, 24b of each CPU. Since the interface units 24a, 24b operate in lock-step synchronism checking can be performed by monitoring the operating states of the paired interface units 24a, 24b by a continuous comparison of certain of their internal states. This approach is implemented by using one stage of a state machine (not shown) contained in the unit 24a of CPU 12A, and comparing each state assumed by that stage with its identical state machine stage in the interface unit 24b. All units of the interface units 24 use state machines to control their operations. Preferably, therefore, a state machine of the memory interface 70 that controls the data transfers between the interface unit 24 and the MC 26 is used. Thus, a selected stage of the state machine used in the memory interface 70 of the interface unit 24a is selected. An identical stage of a state machine of one of the interface unit 24b is also selected. The two selected stages are communicated between the interface units 24a, 24b and received by a compare circuit contained in both interface units 24a, 24b. As the interface units operate lock-step with one another, the state machines will likewise march through the same identical states, assuming each state at substantially the same moments in time. If an interface unit encounters an error, or fails, that activity will cause the interface units to diverge, and the state machines will assume different states. The time will come when... ..that will bring to the attention of the CPUs 12A (or 12B) that the interface units 24a, 24b of that CPU are no longer in lock-step, and to act accordingly... ..X port, receiving only those message packets transmitted by the router 14A of the sub-processor system 10A (Fig. 1A). The Y port is serviced by the Y interface unit 24b to receive message packets from the router 14B of the companion sub-processor system 10B. However, both interfaces (as well as Mcs 26 and processor 20), as has been indicated, are basically mirror images of one another in that both... ..in both structure and function. For this reason, message packet information, received by one interface unit (e.g., 24a) must be passed for processing also to the companion interface unit (e.g., 24b). Further, since both interface units 24a, 24b will assemble the same message packets for transmission from the X or the Y ports, the message packet being transmitted by the interface unit (e.g., 24b) actually being communicated from the associated port (e.g., the Y port) will also be coupled to the other interface unit (e.g., 24a) for cross-checking for errors. These features are illustrated in Figs. 6... ..receiving portions of the packet receivers 96 (96x, 96y) of the X and Y interface units 24a, 24b are broadly illustrated. As shown, each packet receiver 96x, 96y has a clock... ..receive a corresponding one of the TNet Links 32. The CS FIFOs 102 operate to synchronize the incoming command/data symbols to the local clock of the packet receiver 96, buffering... ..104x, coupled to the MUX 104y of the packet receiver 96y of the Y interface unit 24b by the cross-link connection 36(sub(x)). In similar fashion, information received at the Y port is coupled to the X interface unit 24a by the cross-link connection 36(sub(y)). In this manner, the command/data... ..packets received at one of the X,

Y ports by the corresponding X, Y, interface unit 24a, 24b is passed to the other so that both will process and communicate the same information on to other components of the interface units 24 and/or memory 28.

Continuing with Fig. 6, depending upon which port X, Y... ..or the other of the CS FIFOs 102x, 102y for communication to the storage and processing logic 110 of the interface unit 24. The information contained in each 9-bit symbol is an 8-bit byte of... ..the encoding of which is discussed below with respect to Fig. 9. The storage and processing logic 110 will first translate the 9-bit symbols to 8-bit data or command... ..the outputs of the CS FIFOs 102x, 102y are also coupled to a command decode unit in addition to the MUX 104. The command decode unit operates to recognize command symbols (differentiating them from data symbols in a manner that is... ..below), decoding them to generate therefrom command signals that are applied to a receiver control unit, a state machine-based element that functions to control packet receiver operations.

As indicated above... ..at the output of the MUX 104, the receiver control portion of the storage control unit enables CRC check logic 106 to calculate a CRC symbol while the data symbols are... ..below, CS FIFOs are found not only in the packet receivers 96 of the interface units 24, but also at each receiving port of the routers 14 and the I/O... ..an even more important part, and perform a unique function, when a pair of sub-processor systems are operating in duplex mode and the two CPUs 12A and 12B of the sub-processor systems 10A, 10B operate in synchronized, lock-step, executing the same instructions at the same time. When operating in this latter... ..difficult to ensure that the clocking regime of the routers 14A and 14B are exactly synchronized to those of the CPUs 12A and 12B - even when using frequency locked clocking. In... ..used to transmit symbols to a CPU 12 and the clock used by an interface unit 24 to receive those symbols.

The structure of the CS FIFO 102 is diagrammatically illustrated... ..i.e., a packet) or IDLE symbols - except during certain situations (e.g., reset, initialization, synchronization and others discussed below). As explained above, each symbol held in the transmit register 120... ..same symbol leaving the storage queue, allowing each symbol entering the storage queue 126 to settle before it is clocked out and passed to the storage and processing units 110x (and 110y) by the MUX 104x (and 104y). Since the transmit and receive clocks... ..functioning in duplex mode) operate to transmit symbols with near frequency clocking. Even so, clock synchronization FIFOs are used at these other ports to receive symbols transmitted with near frequency clocking, and the structure of these clock synchronization FIFOs are substantially the same as that used in frequency locked environments, i.e., that... ..of the storage queue 126 are nine bits wide; in near frequency environments, the clock synchronization FIFOs use symbol locations of the queue 126 that are 10 bits wide, the extra... ..the faster clock source. To handle this clock drift, the two pointers are effectively re-synchronized periodically.

When the CPUs 12 are paired and operating in duplex mode, all four interface

units 24 operate in lock-step to, among other things, transmit the same data and receive... ..simplex mode, each independent of the other, clocking need only be near frequency.

The interface unit 24 receives a SYNC CLK signal that is used in combination with a SYNC command symbol to initialize and synchronize the Rcv register 124 to the transmitting router 14. When using either near frequency or... ..102X preferably begin from some known state. Incoming symbols are examined by the storage and processing units 110 of the packet receivers 96. The storage and processing units look for, and act upon as appropriate, command symbols. Pertinent here is that when the... ..receives a SYNC command symbol it will be decoded and detected by the storage and processing unit 110. Detection of the SYNC command symbol by the storage and processing unit 110 causes assertion of a RESET signal. The RESET signal, under synchronous control of the SYNC CLK signal, is used to reset the input buffers (including the clock synchronization buffers) to predetermined states, and synchronize them to the routers 14.

The synchronization of the CS FIFOs 102 of the interface units 24 those of one or both routers 14A, 14B is discussed more fully below in the section discussing synchronization.

Packet Transmitter:

Each interface unit 24 is assigned to transmit from and receive at only one of the X or Y ports of the CPU 12. When one of the interface units 24 transmits, the other operates to check the data being transmitted. This is an important... ..shows, in abbreviated form, the packet transmitters 94x, 94y of the X and Y interface units 24a, 24b, respectively. Both packet transmitters are identically constructed, so that discussion of one (packet... ..logic 152 that receives, from the BTE 88 or AVT 90 of the associated interface unit (here, the X interface unit 24a) the data to be transmitted - in doubleword (64-bit) format. The packet assembly logic... ..and Y ports: they are either symbols that make up a message packet in the process of being transmitted, or IDLE symbols, or other command symbols used to perform control functions... ..154, 156. The output of the multiplexer 154 connects to the X port. (The interface unit 24b connects the output of the multiplexer 154 to the Y port.) The multiplexer 156... ..sub(x)) to the checker logic 160 of the packet transmitter 94y (of the interface unit 24b).

A selection (S) input of the multiplexers receives a 1-bit output from an... ..is accessible to the MP 18 via an OLAP (not shown) formed in the interface unit 24, and is written with information that "personalizes," among other things, the interface units 24. Here, the X/Y stage of the configuration register 162 configures the packet transmitter 94x of the X interface unit 24a to communicate the X encoder 150x output to the X port; the output of... ..traffic is present, the operation of the two packet interfaces 94 (and, thereby, the interface units 24 with which they are associated) are continually monitored. Should one of the checkers detect... ..will be asserted, resulting in an internal interrupt being posted for appropriate action by the processors 20.

Message packet traffic operates in the same manner. Assume, for the moment, that the... ..that information, a byte at a time, to the X encoder 150x of both interface units 96, which will translate each byte to encoded 9-bit form. The output of the... ..is checked with that from the packet transmitter 94x. Again, the operation of the interface units 24a, 24b, and the packet transmitters they contain, are inspected for error.

In the same... ..monitored.

Returning for the moment to Fig. 5, if the outgoing message packet is a processor initiated transaction (e.g., a read request), the processors 20 will expect a message packet to be returned in response. Thus, when the BTE... ..will issue a timeout signal to the interrupt logic (Fig. 14A) to thereby notify the processors 20 of the absence of a response to a particular transaction (e.g., a read... ..the access, to name just a few. Also, the area of memory of the memory unit 28 desired to be accessed are identified in the message packets by virtual or I... ..virtual addresses be translated to physical addresses of the memory 28. Finally, interrupts generated by units or elements external to the CPU 12A, are transmitted via message packets to interrupt the processors 20, which are also written to memory 28 when received. All this is handled by the interrupt logic and AVT logic 86, 90.

The AVT logic unit 90 utilizes a table (maintained by the processor 20 in memory 28) containing AVT entries for each possible external source permitted access to the memory 28. Each AVT entry identifies a specific source element or unit and the particular page (a page being nominally 4K (4096) bytes), or portion of a... ..expected" memory accesses. Expected memory accesses are those initiated by the CPU 12 (i.e., processors 20) such as a read request for information from an I/O device. These latter memory accesses are handled by a transaction sequence number (TSN) assigned to each processor initiated request. At about the time the read request is generated, the processors 20 will allocate an area of memory for the data expected to be received in... ..and 26b are, in turn, respectively coupled to the memory interfaces 70 of each interface unit 24a, 24b. The 64-bit doublewords are written to the memory 28 with the upper... ..check bits respectively from the memory interfaces 70 (70a, 70b) of each of the interface units 24a, 24b (Fig. 5).

Referring to Fig. 10, each memory interface 70 receives, from either the bus 82 from the processor interface 60 or the bus 83 from AVT logic 90 (see Fig. 5), of the associated interface unit 24, 64 bits of data to be written to memory. The busses 76 and 83... other for cross-checking between them. Thus, for example, the memory interface 70a (of interface unit 24a) will drive the MC 26a with the "upper" 32 bits of the 64 bits... are check bits, leaving 40 bits unused.

Access Validation:

As previously indicated, components of the processing system 10 external to the CPU 12A (e.g., devices of the I/O packet... not without qualification. Access validation, as implemented by the AVT logic 90 of the interface units 24, operates to prevent the content of the memory 28 from being corrupted ...Accesses to the memory 28 are validated by the AVT logic 90 of each interface unit 24 (Fig. 5), using all of six checks: (1) that the CRC of the message... also are permitted the particular message packet source.

The access validation mechanism of the interface unit 24a, AVT logic 88, is shown in greater detail in Fig. 11. Incoming message packets...and post an interrupt to the interrupt logic 86 (Fig. 5) for action by the processor 20.

The mask operation permits the size of the table of AVT entries to be varied. The content of the AVT mask register 175 is accessible to the processor 20, permitting the processors 20 to optionally select the size of the AVT entry table. A maximum AVT table... 172 allows the AVT size to be matched to the needs of the system. A processing system 10 that includes a larger number of external elements (e.g., the number of... amount of the memory space of memory 28 to the AVT entries. Conversely, a smaller processing system 10, with a smaller number of external elements will not have such a large... set to a logic "ZERO" indicate an nonexistent TNet address, outside the limits of the processing system 10. A received packet with a TNet address outside the allowable TNet range will... in Fig. 11 as being held in the AVT entry register 180 during the validation process. AVT entries have two basic formats: normal and interrupt. The format of a normal AVT...of the AVT input register 170) will result in an error being posted to the processor via an interrupt.

A 12-bit "Permissions" field is included in t AVT entry to... path=0). Denials are logged as interrupts with the interrupt logic, and reported to the processor 20 - if the E field is set to a state ("ONE") that enables error-reporting... e.g., to a "ONE"), the other fields (Upper Bound, etc.) gain new definitions for processing interrupt writes and managing interrupt queues. This is discussed in more detail below in connection... memory 28 will be handled. Set to one state, the requested write operation will be processed normally; set to a second state, write requests specifying addresses with a fractional cache line...be written to a specific queue (interrupt queue) in memory 28, with signalling provided the processors 20 to indicate that a interrupt has been received and "posted," and ready for servicing by the processors 20. Since the interrupt queues are at specific memory locations, the processor can obtain the interrupt data when needed.

An AVT interrupt entry for an interrupt may... by the interrupt logic 86, and extracted from the head of the queue by the processor 20 when servicing the interrupt.

The AVT interrupt entry also includes a 20-bit segment ("Source ID") containing source ID information, identifying the external unit seeking attention by the interrupt process. If the source ID information of the AVT interrupt entry does not match that contained... class" of the interrupt that is used to determine the interrupt level set in the processor 20 (described more fully below); (2) a queue number that is used to select, as...capability to deliver interrupts to a CPU 12 for servicing. For example, an I/O unit may be unable to complete a read or write transaction issued by a CPU because... identify the recipient. These and other errors, exceptions, and irregularities, noted by the I/O units, or the I/O Interface elements, can become the a condition that requires the intervention... the AVT entry register 180 for use by the interrupt logic 86 of the interface unit 24 (Fig. 5), illustrated in greater detail in Fig. 14A.

It is interrupt logic 86...four circular queues specified by the base address information contained in the AVT entry. The processor(s) 20 will then be notified, and it will be up to them as to... ..selected tail queue register 256 by combiner circuit 270, the output of which is the processed by the "mod z" circuit 273 to turn new offset into the queue at which... ..signal. The Queue Full warning signal becomes an "intrinsic" interrupt that is conveyed to the processor units 20 as a warning that if the matter is not promptly handled, later-received interrupt... ..will be discarded.

Incoming message packet interrupts will cause interrupts to be posted to the processor 20 by first setting one of a number of bit positions of an interrupt register 280. Multi-entry queued interrupts are set in interrupt registers 280a for posting to the processor 20; single-entry queue interrupts use interrupt register 280b. Which bit is set depends upon... ..multi-entry queued interrupts, soon after a multi-entry queued interrupt is determined, the interface unit will assert a corresponding interrupt signal (I1) that is applied to decode circuit 283. Decode... ..of register 280a to set, thereby providing advance information concerning the received interrupt to the processor(s) 20, i.e., (1) the type of interrupt posted, and (2) the class of... ..to one another by a compare circuit 279. The update register is writable by the processor 20 to select a register pair for comparison. If the content of the two selected... ..cleared.

Digressing for the moment, there are two basic types of interrupts that concern the processors 20: those interrupts that are communicated to the CPU 12 by message packets, and those...the seven interrupt postings to a latch 288, from which they are coupled to the processor 20 (20a,20b) which has an interrupt register for receiving holding the postings.

In addition... ..change in interrupts (either an interrupt has been serviced, and its posting deleted by the processor 20, or a new interrupt has been posted), a "CHANGE" signal will be issued to the processor interface 60 to inform it that an interrupt posting change has occurred, and that it should communicate the change to the processor 20.

Preferably, the AVT entry register 180 is configured to operate like a single line... ..such as set-associative, fully-associate, or direct-mapped, to name a few.

Coherency:

Data processing systems that use cache memory have long recognized the problem of coherency: making sure that... ..the incoming packet is permitted access are applied to a boundary crossing (Bdry Xing) check unit 219. Boundary check unit 219 also receives an indication of the size of the cache block the CPU 12... ..Len field of the header information from the AVT input register 170. The Bdry Xing unit determines if the data of the incoming packet is not aligned on a cache boundary...time an interrupt will be written to the queued interrupt register 280, to alert the processors 20 that a portion of the incoming data is located in the special queue.

In... ..not, the packet (both header and data) is written to a special queue, and the processors so notified by the intrinsic interrupt process described above. The processors may then move the data from the special queue to cache 22, and later write... ..the cache 22 and the memory 28 is preserved.

Block Transfer Engine (BTE):

Since the processor 20 is inhibited from directly communicating (i.e., sending) information to elements external to the... ..indirect method of information transmission.

The BTE 88 is the mechanism used to implement all processor initiated I/O traffic to transfer blocks of information. The BTE 88 allows creation of... ..BTE registers 300, 302 whose content is coupled to the MUX 306 (of the interface unit 24a; Fig. 5) and used to access the system memory 28 via the memory controllers... ..BTE data structure 304 in the memory 28 of the CPU 12A (Fig. 2). The processors 20 will write a data structure 304 to the memory 28 each time information is... ..begin on a quadword boundary, and the BTE registers 300, 302 are writable by the processors 20 only. When a processor does write one of the BTE registers 300, 302, it does so with a word...

...the request bit (rc0, rc1) to a clear state, which operates to initiate the BTE process, which is controlled by the BTE state machine 307.

The BTE registers 300, 302 also... ..cause (ec) bit differentiates time-outs and NAKs.

When information is being transferred by the processors 20 to an external unit, the data buffer portion 304b of the data structure 304 holds the information to be transferred. When information from an external unit is received by the processors 20, the data buffer portion 304b is the location targeted to hold the read response information.

The beginning of the data structure 304, portion 304a written by the processor 20, includes an information field (Dest), identifying the external element which will receive the packet... ..the transmitted data is to be written. This information is used by the packet transmitter unit 120 (Fig. 5) to assemble the packet in the form shown in Figs. 3-4...list (el) bit, when set, indicates the end of the chain, and halts the BTE processing.

The interrupt completion (ic) bit, when set, will cause the interface unit 24a to assert an interrupt (BTEComp) which sets a bit in the interrupt register 280... ..the chain pointer).

The interrupt time-out (it) bit, when set, will cause the interface unit 24a to assert an interrupt signal for the processor 20 if the acknowledgement of the access times-out (i.e., if the request timer... ..time), or elicits a NAK response (indicating that the target of the request could not process the request).

Finally, if the check sum (cs) bit is set, the data to be... ..containing the data from which the check sum was formed.

To sum up, when the processors 20 of the CPU 12A desire to send data to an external unit, they will write a data structure 304 to the memory 28, comprising identifier information in portion 304a of the data structure, and the data in the buffer portion 304b. The processors 20 will then determine the priority of the data and will write the BTE register... ..information, and sent.

If the data structure 304 indicates a read request (i.e., the processors 20 are seeking data from an external unit - either an I/O device or a CPU 12), the Len and Local Buffer Ptr... ..receiver 100 (Fig. 5) until the local memory write operation is executed.

Responses to a processor-generated read request to an external unit are not processed by the AVT table logic 146. Rather, when the processors 20 set up the BTE data structure, a transaction sequence number (TSN) is assigned the... ..the BTE 88, which will be an HAC type packet (Fig. 4) discussed above. The processors 20 will also include an memory address in the BTE data structure at which the...302, assume that the foregoing transfer of data from the CPU 12A to an external unit is of a large block of information. Accordingly, a number of data structures would be set up in memory 28 by the processors 20, each (except the last) including a chain pointer to additional data structures, the sum... ..sent. Assume now that a higher priority request is desired to be made by the processors 20. In such a case, the associated data structure 304 for such higher priority request... ..with another BTE operation descriptor.

Memory Controller:

Returning, for the moment, to Fig. 2, interface units 24a, 24b access the memory 28 via a pair of memory controllers (MC) 26a, 26b. The Mcs provide a fail-fast interface between the interface units 24 and the memory 28. The Mcs 26 provide the control logic necessary for accessing... ..in dynamic random access memory (DRAM) logic). The Mcs receive memory requests from the interface units 24, and execute reads and writes as well as providing refresh signals to the DRAMs... ..to provide a 72 bit data path between the memory array 28 and the interface units 24a, 24b, which utilize an SBC-DBD-SbD ECC scheme, where b=4, on a... ..26a, 26b to work together and simultaneously supply a 64-bit word to the interface units 24 with minimum latency, one-half of which (D0) comes from the MC 26a, and the other half (D1) comes from the other MC 26b. The interface unit 24 generate and check the ECC check bits. The ECC scheme used will not only... ..26 bus 25, as well as in internal registers.

From the viewpoint of the interface units 24, the memory 28 is accessed with two instructions: a "read N doubleword" and a... ..doubleword read or a block read format. The signal called "data valid" tells the interface units 24 two cycles ahead of time that read data is being returned or not being returned.

As indicated above, the maintenance processor (MP 18; Fig. 1A) has two means of access to the CPUs 12. One is... ..18 will write a register contained in the OLAP 285 with instructions that permit the processors 20 to build an image of a sequence of instructions in the memory that will permit them (the processors 20) to commence ...to transfer instructions and data from an external (storage) device that will complete the boot process.

The OLAP 285 is also used by the processors 20 to communicate to the MP 18 error indications. For example, if one of the interface units 24 detect a parity error in data received from the memory controller 26, it will... ..and address transfers on the bus 25 between the MC 26a and the corresponding interface unit 24a. The addressing and data transfers on the DRAM data bus, as well as generation... ..the CPU 12.

Packet Routing:

The message packets communicated between the various elements of the processing system 10 (e.g., CPUs 12A, 12B, and devices coupled to the I/O packet... ..First, each TNet Link L connects to an element (e.g., router 14A) of the processing system 10 via a port that has both receive and transmit capability. Each transmit port... ..cycle (i.e., each clock period) of the T(underscore)Clk so that the clock

synchronization FIFO at the receiving end of the transmission will maintain synchronization.

Clock synchronization is dependent upon the mode in which the processing system 10 is operated. If operating in the simplex mode in which the CPUs 12A... ..connect directly to the CPUs may drift with respect to each other. Conversely, when the processing system 10 operates in a duplex mode (e.g., the CPUs operate in synchronized, lock-step operation), the clocks between routers 14 and the CPUs 12 to which they... ..not necessarily phase-locked).

The flow of data packets between the various elements of the processing system 10 is controlled by command symbols, which may appear at any time, even within... ..initiated by a CPU 12, or MP 18, and promulgated to all elements of the processing system 10 by the routers 14 to communicate an event requiring software action by all...command symbol is used in conjunction with near frequency operation as an aid to maintaining synchronization between the two clock signals that (1) transfer each symbol to, and load it in each receiving clock synchronization FIFO, and (2) that retrieves symbols from the FIFO.

SLEEP: This command symbol is sent by any element of the processing system 10 to indicate that no additional packet (after the one currently being transmitted, if... ..received.

SOFT RESET (SRST): The SRST command symbol is used as a trigger during the processes ("synchronization" and "reintegration," described below) that are used to synchronize symbol transfers between the CPUs 12 and the routers 14A, 14B, and then to place... ..SYNC command symbol is sent by a router 14 to the CPU 12 of the processing system 10 (i.e., the sub-processor systems 10A/10B) to establish frequency-lock synchronization between CPUs 12 and routers 14A, 14B prior to entering duplex mode, or when in duplex mode to request synchronization, as will be discussed more fully below. The SYNC command symbol is used in conjunction... ..or duplex to simplex), among other things, as discussed further below in the section on Synchronization and Reintegration.

THIS LINK BAD (TLB): When any system element receiving a symbol from a TNet link L (e.g., a router, a CPU, or an I/O unit) notes an error when receiving a command symbol or packet, it will send a TLB... ..identical pairs of symbols that are compared to one another when pulled from the clock synchronization FIFOs..The DVRG command symbol signals the CPU 12 that a mis-compare has been noted. When received by the CPUs, a divergence detection process is entered whereby a determination is made by the CPUs which CPU may be falling... ..command

symbols described above operate to control message flow between the various elements of the processing system 10 (e.g., CPUs 12, router 14, and the like), using principally the BUSY and "end node" (i.e., a CPU 12 or I/O unit 17 - Fig. 1) may not assert backpressure because one of its transmit ports is backpressured... Improperly addressed packets are discarded by the router 14.

When a system element of the processing system 10 receives a BUSY command symbol on a TNet link L on which it... other command symbols (READY, BUSY, etc.).

Whenever a TNet port of an element of the processing system 10 detects receipt of a READY command symbol, it will terminate transmission of FILL... receives.

As will be seen, all elements (e.g., router 14, CPUs 12) of the processing system 10 that connect to a TNet link L for receiving transmitted symbols will receive those symbols via a clock synchronization (CS) FIFO. For example, as discussed above, the interface units 24 of CPUs 12 include all CS FIFOs 102x, 102y (illustrated in Fig. 6). The... depth to allow for speed matching, and the elastic FIFOs must provide sufficient depth for processing delays that may occur between transmission of a BUSY command symbol during receipt of another data byte in packet B. As packet A progresses to the next router, the process would be repeated. If the router 14 displaces more data bytes than the FIFO can... irrespective of its own findings.

SLEEP Protocol:

The SLEEP protocol is initiated by a maintenance processor via a maintenance interface (an on-line access port - OLAP), described below. The SLEEP protocol... reintegrate a slice of the system 10. Routers 14 must be idle (no packets in process) in order to change modes without causing data loss or corruption. When a SLEEP command symbol is received, the receiving element of processing system 10 inhibits initiation of transmission of any new packet on the associated transmit port... The HALT command symbol provides a mechanism for quickly informing all CPUs 12 in a processing system 10 that is necessary to terminate I/O activity (i.e., message transmissions between... CPUs that receive HALT command symbols on either of their receive ports (of the interface units 24) will post an interrupt to the interrupt register 280 if the system halt interrupt... interrupt; Fig. 14A).

The CPUs 12 may be provided with the ability to disable HALT processing. Thus, for example, the configuration registers 75 of the interface units 24 can include a "halt enable register" that, when set to a predetermined state (e.g., ZERO) disables HALT processing, but reporting detection of a HALT symbol as an error.

Router Architecture:

Referring now to... simplified block diagram of the router 14A is illustrated. The other routers 14 of the processing system 10 (e.g., routers 14B, 14', etc.) are of substantially identical construction and, therefore... these ports 4, 5 are structured to operate in a frequency locked environment when a processing system 10 is set for duplex mode operation. In addition, when in duplex mode, a... 5)) will receive the command/data symbols from the CPUs, pass them through the clock synchronization FIFOs 518 (discussed further below), and compare each symbol exiting the clock synchronization FIFOs with a gated compare circuit 517. When duplex operation is entered, a configuration register... 517 to activate the symbol by symbol comparison of the symbols emanating from the two synchronization FIFOs 518 of the router input logic 502 for the ports 4 and 5. Of... to that received, at substantially the same time, by the other port input.

To maintain synchronization in the duplex mode, the two port outputs of the router 14A that transmit to... mode, are duplicated by the routers 14, and returned to both CPUs.) The output logic units 504(sub(4)), 504(sub(5)) that are coupled directly to the CPUs 12 will... message packet identifies only one of the duplexed CPUs 12, e.g., CPU 12A) in synchronized fashion, presenting those symbols in substantially simultaneous fashion to the two CPUs 12. Of course, the CPUs 12 (more accurately, the associated interface units 24) receive the transmitted symbols with synchronizing FIFOs of substantially the same structure as that illustrated in Fig. 7A so that, even... from the

FIFO structures by both CPUs 12 on the same instruction cycle, maintaining the synchronized, lock-step operation of the CPUs 12 required by the duplex operating mode.

As will... ..in conjunction with configuration data written to registers contained in control logic 509 by the maintenance processor 18 (via the on-line access port 285' and serial bus 19A; see Fig. 1A...links L. The input logic 505 of each port input 502 also assists in maintaining synchronization - at least for those ports sending symbols in the near-frequency environment - by removing received... ..slower-receiving element receiving symbols from a faster-sending element could overload the input clock synchronization FIFO of the slower-receiving element. That is, if a slower clock is used to pull symbols from the clock synchronization FIFO put there by a faster clock, ultimately the clock synchronization FIFO will overflow.

The preferred technique employed here is to periodically insert SKIP symbols in... ..stream to avoid, or at least minimize, the possibility of an overflow of the clock synchronization FIFO (i.e., clock synchronization FIFO 518; Fig. 20A) of a router 14 (or CPU 12) due to a T... ..being slightly higher in frequency than the local clock used to pull symbols from the synchronization FIFO. Using SKIP symbols to by-pass a push (onto the FIFO) operation has the... ..stall each time a SKIP command symbol is received so that, insofar as the clock synchronization FIFO is concerned, the transmitting clock that accompanied the SKIP symbol was missing.

Thus, logic... ..the port inputs 502 will recognize, and key off receipt of, SKIP command symbols for synchronization in the near frequency clocking environment so that nothing is pushed onto the FIFO, but... ..14, or between routers 14, or between a router 14 and an I/O interface unit 16A - Fig. 1) at a 50 Mhz rate, this allows for a worst case frequency... ..symbol by supplying FILL or IDLE symbols (which are received and pushed onto the clock synchronization FIFOs, but are not passed to the elastic FIFOs). In short, each elastic FIFO 506...received symbols are then communicated from the input register 516 and applied to a clock synchronization FIFO 518, also by the T(underscore)Clk. The clock synchronization FIFO 518 is logically the same as that illustrated in Figs. 8A and 8B, used in the interface units 24 of the CPUs 12. Here, as Fig. 20A shows, the clock synchronization FIFO 518 comprises a plurality of registers 520 that receive, in parallel, the output of... ..516. Associated with each of the registers 520 is a two-stage validity (V) bit synchronizer 522, shown in greater detail in Fig. 20B, and discussed below. The content of each... ..registers 520, together with the one-bit content of each associated two-stage validity bit synchronizer 522, are applied to a multiplexer 524, and the selected register/synchronizer pulled from the FIFO, and coupled to the elastic FIFO 506 by a pair of... ..is determined the state of the Push Select signal provided by a push pointer logic

unit 530; and, selection of which register 520 will supply its content, via the MUX 524... ..and loading of the register 520 selected by the push pointer logic 530. Similarly, the synchronization FIFO control logic 534 receives the clock signal local to the router (Rcv Clk) to... ..pointer logic 532.

Digressing for a moment, and referring to Fig. 20B, the validity bit synchronizer 522 is shown in greater detail as including a D-type flip-flop 541 with... ..530 (Fig. 20A) selects the register 520 of the FIFO with which the validity bit synchronizer is associated for receipt of the next symbol - if not a SKIP symbol.

The delay... ..Truth Table, below). The D-type flip-flop 543 acts as an additional stage of synchronization, ensuring a stable level at the V output relative to the local Rec Clk. The... ..flip-flop 542, allowing the Pull signal (a periodic pulse from the sync FIFO Control unit 534) to clear the validity bit on this validity synchronizer 522 when the associated register 520 has been read. (Table omitted)

In summary, the validity synchronizer 522 operates to assert a "valid" (V) signal when a symbol is loaded in a...blocked from being routed out a particular port because another message is already in the process of being routed out that port. However, that other message in turn is also blocked...an incoming message packet bound for the CPUs

will be replicated by the crossbar logic unit by routing the message packet to both port output 504(sub(4)) and 504(sub... ..P) identifies which of path (X or Y) should be used for accessing two sub-processing the device.

The routers 14 provide a capability of constructing a large, versatile routing network for, for example, massively parallel processing architectures. Routers are configured according to their location (i.e., level) in the network by...j)) and 509(sub(k)) are such that bits "def" are used in the algorithmic process, then bits "abc" of the Region ID are compared to the content of the Device... ..the route to default register 509(sub(f))) to the final stage of the selection process: check logic 602. Check logic 602 operates to check the status of the port output...a lower level router, and may be located in one or another of the sub-processing systems 10A, 10B. Whether a router is an upper level or lower level router depends... ..of CPUs 12 and I/O devices 16 to one another, forming a massively parallel processing (MPP) system. Other such MPP systems may exist, and it is those routers configured as... ..captured. As soon as the message packet's Destination ID is so captured, the selection process begins, proceeding to the development of a target port address that will be used to...an error that will be posted to the MP18 via the router's (or interface unit's) OLAP for action.

Digressing, it should be appreciated that these protocol rules observed by the routers 14 are also observed by the CPUs 12 (i.e., interface units 24) and I/O packet interfaces 17.

Finally, when the router 14A is in the... ..directly with the CPUs 12A, 12B, and duplex mode is used, a duplex operation logic unit 638 is utilized to coordinate the port output connected to one of the CPUs 12A... ..was able to write instructions to the OLAP 285 that would be executed by the processors 20 to build a small memory image and routine to permit the CPU 12 to... ..the clock generation circuit design. There will be one clock generator circuit in each sub-processor system 10A/10B (Fig. 1) to maintain synchronism. Designated generally with the reference numeral 650... ..used by the various elements (e.g. CPU. 12, routers 14, etc.) of the sub-processor system containing the clock circuit 650 (e.g., 10A).

The clock generator 654 is shown...The 50 Mhz clock signals produced by the counter 663 are distributed throughout the sub-processor system where needed.

Turning now to Fig. 25, there is illustrated the interconnection and use... ..the clock circuits 650 used to develop synchronous clock signals for a pair of sub-processor systems 10A, 10B (Fig. 1) for frequency locked operation. As illustrated in Fig. 25, the two CPUs 12A and 12B of the sub-processor systems 10A, 10B each have a clock circuit 650, shown in Fig. 25 as clock... ..654B of both CPUs 12. A driver and signal line 667 interconnects the two sub-processor systems to deliver the M(underscore)CLK signal developed by the oscillator circuit 652A to the clock generator 654B of the sub-processor system 10B. For fault isolation, and to maintain signal quality, the M(underscore)CLK signal is delivered to the clock generator 654A of the sub-processor system 10A through a separate driver and a loopback connection 668. The reason for the... ..the cable (not shown) will establish the connection shown in Fig. 25 between the sub-processor systems 10A, 10B; connected another way, the connections will be similar, but the oscillator 652B... ..Fig. 25, the M(underscore)CLK signal produced by the oscillator circuit 652A of sub-processing system 10A is used by both sub-processing systems 10A, 10B as their respective SYNC CLK signals and the various other clock signals... ..produced by the clock generators 654A, 654B. Thereby, the clock signals of the paired sub-processing systems 10A, 10B are synchronized for the frequency locked operation necessary for duplex mode.

The VCXOs 662 of the clock... ..This allows both clock generators 654A, 654B to continue to provide to the two sub-processing systems 10A, 10B clock signals in the face of improper operation of the oscillator circuit 652A, although the sub-processor systems may no longer be frequency-locked.

The LOCK signals asserted by the phase comparators... ..LOCK signal signifies that the 50 Mhz signals produced by a clock generator 654 are synchronized, both in phase and in frequency, to the M(underscore)CLK signal. Thus, if either... ..signal that accompanies the symbol stream, and is used to push symbols onto the clock synchronizing FIFO of the receiving element (router 14, or CPU 12) is substantially identical in frequency... ..not phase, to that of

the receiving element used to pull symbols from the clock synchronization FIFOs. For example, referring to Fig. 23, which illustrates symbols being sent from the router... clock (Local Clk). The former (Rcv Clk) is used to push symbols onto the clock synchronization FIFOs 126 of each CPU, whereas the latter is used to pull symbols from the... much higher frequency clock signal. In such situations provision must be made to ensure that synchronization is maintained between the two CPUs as to symbols pulled from the clock synchronization FIFOs 126 of each.

Here, a constant ratio clocking mechanism is used to control operation of the two clock synchronization FIFOs 126, providing the clock signal that pulls symbols from the two FIFOs at the... control mechanism is shown, designated with the reference numeral 70. As Fig. 26A illustrates, clock synchronization FIFO control mechanism 700 includes an pre-settable, multi-stage serial shift register 702, the... ratio of the clock signal at which symbols are communicated and pushed onto the clock synchronization FIFOs 126 to the frequency of the clock signal used locally. Here, a 15 stages... that will be used as the Local Clk signal to pull symbols from the clock synchronization FIFOs 126, and to operate (update) the pull pointer counter 130. The selected output is... of the CPU 12 to the clock signal used to push symbols onto the clock synchronization FIFO 126, Rcv Clk, the serial shift register is preset so that M stages of... duplexed CPUs 12 with a 50 Mhz clock. Thus, symbols are pushed onto the clock synchronization FIFOs 126 of the CPUs at a 50 Mhz rate. Assume further that the clock... of the MUX 704, which produces the clock signal that pulls symbols from the clock synchronization FIFOs 126, Rcv Clk, will contain, for each 100 ns period, five clock pulses. Thus... five symbols will be pushed onto, and five symbols will be pulled from, the clock synchronization FIFOs 126.

This example is symbolically shown in Fig. 26B, while the timing diagram shown... labelled "IN" in Fig. 27) of the Rcv Clk will push symbols onto the clock synchronization FIFOs 126. During that same 100 ns period, the serial shift register 702 circulates a... clocks which would require additional storage (i.e., an increase in the size of the synchronization FIFO) and impose more latency.

The constant ratio clock circuit presented here (Figs. 26) is... frequency to a clock regime of a different, higher frequency. The use of a clock synchronization FIFO is necessary here for compensating effects of signal delays when operating in synchronized, duplexed mode to receive pairs of identical command/data symbols from two different sources. However... so long as there are at least two registers in the place of the clock synchronization FIFO. Transferring data from a higher-frequency clock regime to a lower frequency clock regime... a wide range of possible clock ratios.

I/O Packet Interface:

Each of the sub-processor systems 10A, 10B, etc. will have some input/output capability, implemented with various peripheral units, although it is conceivable that the I/O of other sub-processor systems would be available so that a sub-processing system may not necessarily have local I/O. In any event, if local I/O... device (e.g., a signal line) would be received by the I/O packet interface unit 16 and used to form an interrupt packet that is sent to the CPU 12... OLAP bus, configuration information.

On-Line Access Port:

The MP 18 connects to the interface unit 24, memory controller (MC) 26, routers 14, and I/O packet interfaces with interface signals... OLAP 258 is essentially the same, regardless of what element (e.g. router 14, interface unit 24, etc.) it is used with. Fig. 28 diagrammatically illustrates the general structure of the... circuit chip used to implement certain of the elements discussed herein. For example, each interface

unit 24, memory controller 26, and router 14 is implemented by an application specific integrated circuit... of the OLAP 158 shown in Fig. 28 describes the OLAP associated with the interface unit 24, the MC 26, and the router 14 of the system.

As Fig. 28 shows...asymmetric variables, a "soft-vote" (SV) logic element 900 (Fig. 30A) is provided each interface unit 24 of each CPU 12. As Fig. 30 illustrates, the SV logic elements 900 of each interface unit 24 are connected to one another by a 2-bit SV bus 902, comprising bus lines 902a and 902b. Bus lines 902a carry one-bit values from the interface units 24 of CPU 12A to those of CPU 12B. Conversely, bus line 902b carries one... ..the CPU 12A.

Illustrated in Fig. 30B, is the SV logic element 900a of interface unit 24a of CPU 12A. Each SV logic element 900 is substantially identical in construction and... ..900a should be understood as applying equally to the other logic elements 900a (of interface unit 24b, CPU 12A), and 900b (of the interface units 24a, 24b of CPU 12B) unless noted otherwise. As Fig. 30B illustrates, the SV logic interface units 24a, 24b of the CPU 12A can communicate asymmetrical variables to each other.

In a... ..to the remote register 907 of logic element 902a (and that of the other interface unit 24b).

The logic elements 902 form a part of the configuration registers 74 (Fig. 5). Thus, they may be written by the processor unit(s) 20 by communicating the necessary data/address information over at least a portion of... ..local and remote registers 906 and 907.

The MUX 914 operates to provide each interface unit 24 of CPU 12A with selective use of the bus line 902a for the SV logic elements 900a, or for communicating a BUS ERROR signal if encountered during the reintegration process (described below) used to bring a pair of CPUs 12 into lock-step, duplex operation... ..same time, write the enable registers 912 of the logic element 900 of both interface units 24 of each CPU. One of the two logic elements 900 of each CPU will... ..it is the output enable registers 912 associated with the logic elements 900 of interface units 24a of both CPUs 12A, 12B that are written to enable the associated drivers 916. Thus, the output registers 904 of the interface units 24a of each CPU will be communicated to the bus lines 902; that is, the... ..to the bus line 902a, while the output register associated with logic element 900b interface unit 24a of CPU 12B is communicated to bus line 902b. The CPUs 12 will both... ..again written by each CPU, followed again by reading the remote input registers 907. This process is repeated, one bit at a time, until the entire variable is communicated from the... ..each CPU 12 to the remote input register of the other. Note that both interface units 24 of CPU 12B will receive the bit of asymmetric information.

One example of use... ..elements 900 are also used to communicate bus errors that may occur during the reintegration process to be described. When reintegration is being conducted, a REINT signal will be asserted. As... ..ERROR signal is selected by the MUX 914 and communicated to the bus line 902a.

Synchronization:

Proper operation of the sub-processing systems 10A, 10B (Figs. 1A, 2) whether operating independently (simplex mode), or paired and operating in synchronized lock-step (duplex mode), requires assurance that data communicated between the CPUs 12A, 12B and... ..routers 14A, 14B will be received properly, and that any initial content of the clock synchronization FIFOs 102 (of CPUs 12A, 12B; Fig. 5) and 519 (of routers 14A, 14B; Fig... ..erroneously interpreted as data or commands. The push and pull pointers of the various clock synchronization FIFOs 102 (in the CPUs 12) and 518 (in the routers 14) need to be... ..apart, and presetting the associated FIFO queues to some known state. This done, all clock synchronization FIFOs are initialized for near ...in order to properly implement the lock-step operation of duplex mode operation, the clock synchronization FIFOs must be synchronized to operate with the particular source from which they receive data in order accommodate any... ..14A, 14B to the CPUs 12A, 12B must be accounted for. It is the clock synchronization FIFOs 102 of the paired CPUs 12 that operate to receive message packet symbols, adjust... ..and present symbols to the two CPUs in a simultaneous manner to maintain lock-step synchronization necessary for duplex mode operation.

In similar fashion, each symbol received by the routers 14A... ..the CPUs (which is discussed further hereinafter). Again, it is the function of the clock synchronization FIFOs 518 of the routers 14A, 14B that receive message

packets from the CPUs 12... ..so that the symbols received from the two CPUs 12 are retrieved from the clock synchronization FIFOs simultaneously.

Before discussing how the clock synchronization FIFOs of the CPUs and routers are reset, initialized, and synchronized, an understanding of their operation to maintain synchronous lock-step duplex mode operation is believed helpful. Thus, referring for the moment to Fig. 23, the clock synchronization FIFOs 102 of the CPUs 12A, 12B that receive data, for example, from the router... ..underscore)Clk, from the router 14A to the CPU 12B.

Consider operation of the clock synchronization FIFOs 102(sub(x)), 102(sub(y)), to receive identical symbol streams during duplex operation... ..held by the push and pull pointer counters 128, 130 for the CPU 12A (interface unit 24a), and the content of each of the four storage locations (byte 0. ..., byte 3... ..6 show the same thing for the FIFO 102(sub(y)) of CPU 12B interface unit 24a for each symbol of the duplicated symbol stream.

Assuming the delay 640 is no...0" locations of the queues 126. This is because (1) the FIFOs 102 have been synchronized to operate in synchronism (a process described below), and (2) the push pointer counters 128 are clocked by the clock signal... ..of the symbol stream transmitted by the router 14A will be pulled from the clock synchronization FIFOs 102 of the CPUs 12A, 12B simultaneously, maintaining the required synchronization of received data when operating in duplex mode. In effect, the depths of the queues... ..order to achieve the operation just described with reference to Table 6, the reset and synchronization process shown in Fig 31A is used. The process not only initializes the clock synchronization FIFOs 102 of the CPUs 12A, 12B for duplex mode operation, but also operates to adjust the clock synchronization FIFOs 518 (Fig. 19A) of the CPU ports of each of the routers 14A, 14B for duplex operation. The reset and synchronization process uses the SYNC command symbol to initiate a time period, delineated by the SYNC CLK signal 970 (Fig. 31B), to reset and initialize the respective clock synchronization FIFOs of the CPUs 12A and 12B and routers 14A, 14B. (The SYNC CLK signal... ..It is of a lower frequency than that used to receive symbols by the clock synchronization FIFOs, T(underscore)Clk. For example, where T(underscore)Clk is approximately 50 MHz, the... ..signal is approximately 3.125 MHz.)

Turning now to Fig. 31A, the reset and initialization process begins at step 950 by switching the clock signals used by the CPUs 12A, 12B and routers 14A, 14B as the transmit (T(underscore)Clk) and the unit's local clock (Local Clk) clock signals so that they are derived from the same... ..In addition, configuration registers in the CPUs 12A, 12B (configuration registers 74 in the interface units 24) and the routers 14A, 14B (contained in control logic unit 509 of routers 14A, 14B) are set to the FreqLock state.

The following discussion involves step 952, and makes reference to the interface unit 24 (Fig.5), router 14A (Fig. 19A) and Figs. 31A and 31B. With the clock... ..otherwise be sent followed by a self-addressed message packet. Any message packet in the process of being received and retransmitted when the SLEEP command symbols are received and recognized by... ..per the destination address). The SLEEP command symbol operates to "quiece" router 14A for the synchronization process. The self-addressed message packet sent by the CPU 12A, when received back by the... ..message packet sent after the SLEEP command symbol would necessarily have to be the last processed by the router 14A.

At step 954 the CPU 12A checks to see if it...the router will assert a RESET signal 972 that is applied to the two clock synchronization FIFOs 518 contained in the input logic 505(sub(4)), 505(sub(5)) of the... ..receive symbols directly from CPUs 12A, 12B. RESET, while asserted, will hold the two clock synchronization FIFOs 518 in a temporarily non-operating reset state with the push and pull pointer... ..As each of the CPUs 12 receive SYNC symbols are detected by the storage and processing units of the packet receivers 96 (Figs. 5 an 6) cause the RESET signal to be asserted by the packet receivers 96 (actually, storage and processing elements 110; Fig. 6) of each CPU 12. the RESET signal is applied to the... ..4))), CPUs 12 and routers 14A, 14B de-assert the RESET signals, and the clock synchronization FIFOs of the CPUs 12A, 12, and routers 14A, 14B are released from their reset... ..the delay, the router 14A and CPUs 12 resume pulling data from their respective clock synchronization FIFOs and resume normal operation. The clock synchronization FIFOs of the router 14A begin pulling symbols from the queue

(previously set by RESET... from the CPU 12A with the T(underscore)Clk will be pushed onto the clock synchronization FIFO at, for example, queue location 0 (or whatever other location pointed to by the... 0 (or whatever other location the push pointer was set to by RESET). The clock synchronization FIFOs of the router 14A are now synchronized to accommodate whatever delay 640 may be present in one communications path, relative to the... and the CPUs 12A, 12B.

Similarly, at the same virtual time, operation of the clock synchronization FIFOs 102 of both CPUs 12A, 12B is resumed, synchronizing them to the router 14A. Also, the CPUs 12A, 12B quit sending the SLEEP command... in favor of READY symbols, and resume message packet transmission, as appropriate.

That completes the synchronization process for the router 14A. However, the process must also be performed for the router 14B. Thus, the CPU 12A returns to step... however, assuming that the CPUs 12A, 12B are operating in duplex mode, the method and apparatus used to detect and handle a possible error, resulting in divergence of the CPUs from... via a message packet destined for a peripheral device of one or the other sub-processor systems 10A, 10B. Depending upon the destination of the outgoing message packet, step 1002 will... router 14 will issue an ERROR signal to the router control logic 509, causing the process to move to step 1004 where the router 14 detecting divergence will transmit a DVRG... time outs to occur. A router detecting divergence (without also detecting any simple link error) buys itself time to check the CRC of the received message packet by waiting for the... router 14, or received, all further message packets received from the CPUs and in the process of being routed when divergence was detected, or the DVRG symbol received, will be passed... 1010) contained in a one of the configuration registers 74 (Fig. 5) of the interface unit 24 of each CPU.

Returning for the moment to step 1006, the determination of which... local" is meant to refer to the router 14A, 14B contained in the same sub-processor system 10A, 10B as the CPU. For example, referring to Fig. 1A, router 14A is... bit mentioned above: the bit contained in one of the configuration registers 74 of interface unit 24 (Fig. 5) of each CPU. When set to a first state, that particular CPU... the other CPU. In response, the state machines (not shown) within the control and status unit 509 (Fig. 19A) changes the "favorite" bits described above.

A few examples may facilitate understanding... DVRG symbol will echo that symbol to the routers 14A, 14B, start its internal divergence process timer, and begin determination of whether to continue or terminate. Having received a TLB symbol... to diverge with no errors reported. This can happen only if software (running on the processors 20) uses known divergent data to alter state. For example, suppose each CPU 12 has... number of the CPU 12A will differ from that of the CPU 12B. If the processors use the serial number to change the sequence of instructions executed (say, by branching if the serial number comes after some value) or to modify the value contained in a processor register, the complete "state" of the CPUs 12 will differ. In such cases, the "asymmetrical... of the primary CPU simply allows one CPU, and thereby the system 10, to continue processing without software intervention.

- An error at the output of the interface unit 24 of a CPU 12 will be detected by the router 14A, 14B, depending upon... router 14A, 14B that connects to a CPU 12 will be detected by the interface unit 24 of the affected CPU. The CPU will send a TLB symbol to the faulty... possible failure and, without external intervention, and transparently to the system user, remove the failing unit (CPU 12A or 12B, or router 14A or 14B) from the system to obviate or... reintegration." The discussion will refer to the CPUs 12A, 12B; routers 14A, 14B, and maintenance processor 18A, 18B shown forming parts of the processing system 10 illustrated in Fig. 1A. In addition, discussion will refer to the processors 20a, 20b, the interface units 24a, 24b, and the memory controllers 26a, 26b (Fig. 2) of the CPUs 12A, 12B as single units, since that is the way they function.

Reintegration is used to place two CPUs in... both of the paired CPUs at virtually the same time.

The major steps in the process for changing from simplex mode operation of the one on-line CPU to duplex mode...
...greater detail by the flow diagrams of Figs. 33A - 33D, generally are:

1. Setup and synchronize the two CPUs (one on-line, the other off-line) and their connected routers to... ..the memory of the on-line CPU to the off-line CPU, maintaining a tracking process that monitors changes in the memory of the on-line CPU that have not been... ..and may need to be copied over to, the off-line CPU;
3. Setup and synchronize the CPUs to run a delayed (slave) duplex mode from the same instruction stream (lock... ..will write the predetermined registers (not shown) of the control registers 74 in the interface units 24 of CPUs 12A and 12B, to a next state (after a soft operation) in... ..the off-line CPU 12B.

Next, a sequence is entered (steps 1060 - 1070) that will synchronize the clock synchronization FIFOs of the CPUs 12A, 12B and routers 14A, 14B in much the same fashion... ..the same steps described above in connection with the discussion of Figs. 31A, 31B to synchronize the clock synchronization FIFOs. The on-line CPU 12A will send the sequence of a SLEEP symbol, self-addressed message packet, and SYNC symbol which, with the SYNC CLK signal, operates to synchronize CPUs and routers. Once so synchronized, the on-line CPU 12A then, at step 1066, sends a Soft Reset (SRST) command... ..of all configuration registers and control registers (e.g., configuration registers 74 of the interface units 24) cache, and the like to memory 28 of the ...time to have the system 10 off-line for reintegration. For that reason, the reintegration process is performed in a manner that allows the on-line CPU to continue executing user... ..not match that of the off-line CPU. The reason for this is that normal processing by the processor 20 of the on-line CPU can change memory content after it has been copied... ..when a memory location is written in the on-line CPU 12A during the reintegration process it is marked as "dirty;" second, all copying of memory to the off-line CPU... ..may, however, limit the ability to detect two-bit errors. But, since the memory copying process will last for a only relatively short period of time, this risk is believed acceptable... ..memory location in CPU 12A is made (either an incoming I/O write, or a processor write operation). The returning data (that was copied over to the off-line CPU) would... ..controller 26 (Fig. 2) of the on-line CPU to monitor memory locations in the process of being copied over to the off-line CPU 12B. The memory controller uses a...within the block had been written by another operation (e.g., a write by the processor 20, an I/O write, etc.), that prior write operation will flag the location in... ..still must be copied over to the off-line CPU 12B.

Returning to the reintegration process, and now to Fig. 33B, the memory tracking (AtomicWrite mechanism and using ECC to mark... ..entails writing a reintegration register (not shown; one of the configuration registers 74 of interface unit 24 - Fig. 5) to cause a reintegration (REINT) signal to be asserted. The REINT signal... ..is left alone. Throughout the incremental copy operations, the normal actions of the on-line processor will mark some memory locations dirty.

Several passes of incremental copying will need to be... ..the number of successful WriteConditional operations at the end of each pass through memory, the processors 20 can determine the effect of a given pass compared to the previous pass. When the benefits drop off, the processors 20 will give up on the precopy operations. At this point the reintegration process is ready to place the two CPUs 12A, 12B into lock-step operation.

Thus, the... ..in Fig. 33C, where at step 1100, the on-line CPU 12A momentarily halts foreground processing, i.e., execution of a user application. The remaining state (e. g., configuration registers, cache, etc.) of the on-line processors 20 and its caches is then read and written to a buffer (series of memory... ..to the off-line CPU 12B, together with a "reset vector" that will direct the processor units 20 of both CPUs 12A, 12B to a reset instruction.

Next, step 1106 will quiesce... ..to ensure that the FIFOs of the routers are clear, that the FIFOs of the processor interfaces 24 are clear, and no further incoming I/O message packets are forthcoming. At... ..symbol will be received and acted upon by both CPUs 12A, 12B, to cause the processor units 20 of each CPU to jump to the location in memory 28 containing the reset... ..a subroutine that will restore the stored state of both CPUs 12A, 12B to the processor units 20, caches 22, registers, etc. The CPUs 12A, 12B will then begin executing the same... ..enabling of

the ECC bit to mark dirty locations must now be disabled, since the processors are doing the same thing to the same memory. During this stage of the reintegration... ..encountered by CPU 12A.

Meanwhile, the bus error in the CPU 12A will cause the processor unit 20 to be forced into an error-handling routine to determine (1) the cause of... ..error was caused by an attempt to read a memory location marked dirty. Accordingly, the processor unit 20 will initiate (via the BTE 88 -- Fig. 5) the AtomicWrite mechanism to copy the...the SRST symbols are now received by the CPUs 12A, 12B, they will cause both processor units 20 of the CPUs to be reset to start from the same location with the... ..will periodically update, e.g., a database or audit file that is indicative of the processing of the primary CPU up to that point in time of the update. Should the... ..in error-checking redundancy to the CPU 12B, in the same manner that the individual processor units 20a, 20b of the CPU 12A provide fail-fast, fault tolerance for the CPU - when... ..cost system is applicable, as illustrated in Fig. 34. As shown in Fig. 34, a processing system 10' includes the CPU 12A and routers 14A, 14B structured as described above. The... ..and the CPUs are also the same.

Thus, the CPU 12B' comprises only a single processor unit 20' and associated support components, including the cache 22', interface unit (IU) 24', memory controller 26', and memory 28'. Thus, while the CPU 12A is structured in the manner shown in Fig. 2, with cache processor unit, interface unit, and memory control redundancies, approximately one-half of those components are needed to implement CPU... ..stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... ..i.e., parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of... ..inputs connected to receive the data output

from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... ..for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the... ..attached, without the assistance of any other CPU 12 in the system. Many prior parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O device becomes... ..rest of the system. Other prior systems provide access to I/O through pairs of processors so that should one of the processors fail, access to the corresponding I/O is still available through the remaining I/O... ..if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... ..multiprocessing system access to every peripheral, as done here, operates to extend the "primary"/"backup" process taught in the above-identified U.S. Patent No. 4,228,496. There, a multiple CPU system may have a primary process may running on one CPU, while a backup process resides in the background on another of the CPUs. Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be... ..e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system... ..CPU can access any I/O without

the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time... indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Claims: ...A2

1. In a processing system including a processor unit and a plurality of peripheral elements coupled to the processor unit for communicating message packets therebetween, a method for verifying a communication path between the processor unit and a one of the plurality of peripheral elements, including the steps of:

the processor unit sending to the one of the plurality of peripheral elements at least one prior message... of the plurality of peripheral elements to response by sending a response message packet;

the processor unit sending to the one of the plurality of peripheral elements a barrier transaction message packet...

6/K/25 (Item 25 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Country	Number	Kind	Date
Type		Pub. Date	Kind
...Compaq Computer Corporation (687797) 20555 SH 249 Houston, Texas 77070-2698 US		19	
Examination...		19	
Available Text		Language	Update
			Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: ...and 08/485,053, filed concurrently herewith.

The present invention is directed generally to data processing systems, and more particularly to a multiple processing system and a reliable system area network that provides connectivity for interprocessor and input/output... and communications systems to general purpose high availability commercial systems. The evolution of fault tolerant computers has been well documented (see D. P. Siewiorek, R. S. Swarz, "The Theory and

Practice... ..and the Jet Propulsion laboratory began to apply fault tolerance to the development of guidance computers for aerospace applications. The 1960's also saw the development of the first AT&T electronic switching systems.

The first commercial fault tolerant machines were introduced by Tandem Computers in the 1970's for use in on-line transaction processing applications (J. Bartlett, "A NonStop Kernel," in proc. Eighth Symposium on Operating System Principles, pp... ..systems were introduced in the 1980's (O. Serlin, "Fault- Tolerant Systems in Commercial Applications," Computer, pp. 19-30, August 1984). Current commercial fault tolerant systems include distributed memory multi-processors, shared-memory transaction based systems, "pair-and-spare" hardware fault tolerant systems (see R. Freiburghouse, "Making Processing Fail-safe," Mini-micro Systems, pp. 255-264, May 1982; U.S. Patent No. 4... ..system.), and triple-modular-redundant systems such as the "Integrity" computing system manufactured by Tandem Computers Incorporated of Cupertino, California, assignee of this application and the invention disclosed herein.

Most applications of commercial fault tolerant computers fall into the category of on-line transaction processing. Financial institutions require high availability for electronic funds transfer, control of automatic teller machines, and... ..telecommunications systems.

Vendors of fault tolerant machines attempt to achieve both increased system availability, continuous processing, and correctness of data even in the presence of faults. Depending upon the particular system architecture, application software ("processes") running on the system either continue to run despite failures, or the processes are automatically restarted from a recent checkpoint when a fault is encountered. Some fault tolerant systems are provided with sufficient component redundancy to be able reconfigure around failed components, but processes running in the failed modules are lost. Vendors of commercial fault tolerant systems have extended fault tolerance beyond the processors and disks. To make large improvements in reliability, all sources of failure must be addressed... ..power supplies, fans and inter-module connections.

The "NonStop," and "Integrity" architectures manufactured by Tandem Computers Incorporated, (both respectively illustrated broadly in U.S. Patent No. 4,228,496 and U... ..assigned to the assignee of this application; NonStop and Integrity are registered trademarks of Tandem Computers Incorporated) represent two current approaches to commercial fault tolerant computing. The NonStop system, as generally... ..above-identified U.S. Patent No. 4,278,496, employs an architecture that uses multiple processor systems designed to continue operation despite the failure of any single hardware component. In normal operation, each processor system uses its major components independently and concurrently, rather than as "hot backups". The NonStop system architecture may consist of up to 16 processor systems interconnected by a bus for interprocessor communication. Each processor system has its own memory which contains a copy of a message-based operating system. Each processor system controls one or more input/output (I/O) busses. Dual-porting of I/O controllers and devices provides multiple paths to each device. External storage (to the processor system), such as disk storage, may be mirrored to maintain redundant permanent data storage.

This... ..hardware, while fault recovery is the responsibility of the software.

Also, in the Nonstop multi-processor architecture, application software ("process") may run on the system under the operating system as "process-pairs," including a primary process and a backup process. The primary process runs on one of the multiple processors while the backup process runs on a different processor. The backup process is usually dormant, but periodically updates its state in response to checkpoint messages from the primary process. The content of a checkpoint message can take the form of complete state update, or... ..checkpoints were manually inserted in application programs, but currently most application code runs under transaction processing software which provides recovery through a combination of checkpoints and transaction two-phase commit protocols.

Interprocessor message traffic in the Tandem Nonstop architecture includes each processor periodically broadcasting an "I'm Alive" message for receipt by all the processors of the system, including itself, informing the

other processors that the broadcasting processor is still functioning. When a processor fails, that failure will be announced and identified by the absence of the failed processor's periodic "I'm Alive" message. In response, the operating system will direct the appropriate backup processes to begin primary execution from the last checkpoint. New backup processes may be started in another processor, or the process may be run with no backup until the hardware has been repaired. U.S. Patent... ..example of this technique.

Each I/O controller is managed by one of the two processors to which it is attached. Management of the controller is periodically switched between the processors. If the managing processor fails, ownership of the controller is automatically switched to the other processor. If the controller fails, access to the data is maintained through another controller.

In addition to providing hardware fault tolerance, the processor pairs of the above-described architecture provide some measure of software fault tolerance. When a processor fails due to a software error, the backup processor frequently is able to successfully continue processing without encountering the same error. The software environment in the backup processor typically has different queue lengths, table sizes, and process mixes. Since most of the software bugs escaping the software quality assurance tests involve infrequent data dependent boundary conditions, the backup processes often succeed.

In contrast to the above-described architecture, the Integrity system illustrates another approach... ..fault recovery is the logical choice since few modifications to the software are required. The processors and local memories are configured using triple-modular-redundancy (TMR). All processors run the same code stream, but clocking of each module is independent to provide tolerance... ..three streams is asynchronous, and may drift several clock periods apart. The streams are re-synchronized periodically and during access of global memory. Voters on the TMR Controller boards detect and mask failures in a processor module. Memory is partitioned between the local memory on the triplicated processor boards and the global memory on the duplicated TMRC boards. The duplicated portions of the... ..techniques to detect failures. Each global memory is dual ported and is interfaced to the processors as well to the I/O Processors (IOPs). Standard VME peripheral controllers are interfaced to a pair of busses through a Bus... ..the BIMs to switch control of all controllers to the remaining IOP. Mirrored disk storage units may be attached to two different VME controllers.

In the Integrity system all hardware failures... ..reintegrated on-line.

The preceding examples illustrate present approaches to incorporating fault tolerance into data processing systems. Approaches involving software recovery require less redundant hardware, and offer the potential for some... ..have been developed on other systems.

Thus, the systems described above provide fault tolerant data processing either by hardware (e.g, fail-functional, employing redundancy) or by software techniques (fail-fast... ..hardware). However, none of the systems described are believed capable of providing fault tolerant data processing, using both hardware (fail-functional) and software (fail-fast) approaches, by a single data processing system.

Computing systems, such as those described above, are often used for electronic commerce: electronic data interchange (EDI) and global messaging. Today's demands upon such electronic commerce, however, is demanding more and more throughput capacity as the number of users increases and... ..networks such as local area networks (LANS), and the like.

A key requirement for a server architecture is the ability to move massive quantities of data. The server should have high bandwidth that is scalable, so that added throughput capacity can be added... ..response time, latency affects service levels and employee productivity.

The present invention provides a multiple-processor system that combines both of the two above-described approaches to fault tolerant architecture, hardware redundancy and software recovery techniques, in a single system.

Broadly, the present invention includes a processing system composed of multiple sub-processing systems. Each sub-processing system has, as the main processing element, a central processing unit (CPU) that in turn comprises a pair of processors operating in lock-step, synchronized fashion to execute each instruction of an instruction stream at the same time. Each of the sub-processing systems further include an input/output (I/O) system area network system that provides redundant communication paths between various components of the larger

processing system, including a CPU and assorted peripheral devices (e.g., mass storage units, printers, and the like) of a sub-processing system, as well as between the sub-processors that may make up the larger overall processing system. Communication between any component of the processing system (e.g., a CPU and another CPU, or a CPU and any peripheral device, regardless of which sub-processing system it may belong to) is implemented by forming and transmitting packetized messages that are... responsible for choosing the proper or available communication paths from a transmitting component of the processing system to a destination component based upon information contained in the message packet. Thus, the... peripherals, but permits it to also be used for interprocessor communications.

As indicated above, the processing system of the present invention is structured to provide fault-tolerant operation through both "fail... at a variety of points in the various data paths between the (lock-step operated) processor elements of the CPU and its associated memory. In particular, the processing system of the present invention conducts error-checking at an interface, and in a manner... little impact on performance. Prior art systems typically implement error-checking by running pairs of processors, and checking (comparing) the data and instruction flow between the processors and a cache memory. This technique of error-checking tended to add delay to the... error-checking precluded use of off-the-shelf parts that may be available (i.e., processor/cache memory combinations on a single semiconductor chip or module). The present invention performs error-checking of the processors at points that operate at slower rates, such as the main memory and I/O interfaces which operate at slower speeds than the processor-cache interface. In addition, the error-checking is performed at locations that allow detection of errors that may occur in the processors, their cache memory, and the I/O and memory interfaces. This allows simpler designs for... other data integrity checks.

Error-checking of the communication flow between the components of the processing system is achieved by adding a cyclic-redundancy-check (CRC) to the message packets that... "Good" (TPG) or "This Packet Bad" (TPB) - is appended to every packet. A maintenance diagnostic processor can use this information to isolate a link or router element that introduces an error... of topologies, so that alternate paths can be provided between any two elements of a processing system (e.g., between a CPU and an I/O device), for communication in the... so (e.g., by creating a "deadlock" condition, discussed further below).

The CPUs of a processing system are capable of operating in one of two basic modes. a "simplex mode" in... independently of the other, or a "duplex" mode in which pairs of CPUs operate in synchronized, lock-step fashion. Simplex mode operation provides the capability of recovering from faults that are... U.S. Pat. No. 4,228,496 which teaches a multiprocessing system in which each processor has the capability of checking on the operability of its sibling processors, and of taking over the processing of a processor found or believed to have failed). When operating in duplex mode, the paired CPUs both... fault tolerant platform for less robust operating systems (e.g., the UNIX operating system). The processing system of the present invention, with the paired, lock-step CPUs, is structured so that... masked (i.e., operating despite the existence of a fault), primarily through hardware.

When the processing system is operating in duplex mode, each CPU pair uses the I/O system to access any peripheral of the processing system, regardless of which (of the two, or more) sub-processor system the peripheral may be ostensibly a member of. Also, in duplex mode, message packets... message for the CPU pair (from either a peripheral device such as a mass storage unit or from a processing unit), will replicate the message and deliver it to both CPUs of the pair using synchronization methods that ensure that the CPUs remain synchronized. In effect, the

duplex CPU pair, as viewed from the I/O system and other... as a single CPU. Thus, the I/O system, which includes elements from all sub-processing systems, is made to be seen by the duplex CPU pair as one homogeneous system... a multiprocessor system in which the CPU of any one is actually a pair of synchronized, lock-step CPUs.

Yet another important aspect of the present invention is that interrupts issuing... interrupts via the message packet system ensures that they will arrive at duplexed CPUs in synchronized fashion, in the same manner as I/O message packets. Interrupt message packets will contain... the system. In addition, using the same messaging system to communicate data between I/O units and the CPUs and to communicate interrupts to the CPUs preserves the ordering of I... the implementation of a technique of validating access to the memory of any CPU. The processing system, as structured according to the present invention, permits the memory of any CPU to... to handle input/output information transfers between a CPU and any other component of the processor system. Thereby, the individual processor units of the CPU are removed from the more mundane tasks of getting information from memory and out onto the TNet network, or accepting information from the network. The processor unit of the CPU merely sets up data structures in memory containing the data to be... is required, where in memory the response is to be placed when received. When the processor unit completes the task of creating the data structure, the block transfer engine is notified to... response is received, it is routed to the expected memory location identified, and notifies the processor unit that the response was received.

Further aspects and features of the present invention will become... invention, which should be taken in conjunction with the accompanying drawings.

Fig. 1A illustrates a processing system constructed in accordance with the teachings of the present invention, and Figs. 1B and 1C illustrate two alternate configurations of the processing system of Fig. 1A, employing clusters or arrangements of the processing system of Fig. 1A;

Fig. 2 illustrates, in simplified block diagram form, the central processing unit (CPU) that forms a part of each sub-processor system of Figs. 1A - 1C;

Figs. 3A - 3D and 4A - 4C illustrate the construction of... the area network I/O system shown in Fig. 2;

Fig. 5 illustrates the interface unit that forms a part of the CPUs of Fig. 2 to interface the processor and memory with the I/O area network system;

Fig. 6 is a block diagram, illustrating a portion of packet receiver of the interface unit of Fig. 5;

Fig. 7A diagrammatically illustrates the clock synchronization FIFO (CS FIFO) used by the packet receiver section packet receiver shown in Fig. 6;

Fig. 7B is an block diagram of a construction of the clock synchronization FIFO structure shown in Fig. 7A;

Fig. 8 illustrates the cross-connections for error-checking outbound transmissions from the two interface units of a CPU;

Fig. 9 illustrates an encoded (8B to 9B) data/command symbol;

Fig. 10 illustrates the method and structure used by the interface unit of Fig. 5 to cross-check for errors data being transferred to the memory controllers... of a CPU of Fig. 2 to other (external to the CPU) components of the processing system;

Fig. 12 is a block diagram that diagrammatically illustrates the formation of an address... 14A illustrates the logic for posting interrupt requests to queues in memory and to the processor units of the CPU of Fig. 2;

Fig. 14B illustrates the process used to form a memory address for a queue entry;

Fig. 15 is a block... data output constructs formed in the memory of the CPU of Fig. 2 by a processor unit, and containing data to be sent via the area I/O networks shown in Figs. 1A - 1C, and also illustrating the block transfer

engine (BTE) unit of the interface unit of Fig. 5 that operates to access the data output constructs for transmission to the... pair of memory controllers between memory of a CPU of Fig. 2 and its interface unit for accessing from memory 72 bits of data, including two simultaneously-accessed 32-bit words... other for error-checking;

Fig. 19A is a simplified block diagram illustration of the router unit used in the area input/output networks of the processing systems shown in Figs. 1A - 1C;

Fig. 19B illustrates comparison on two port inputs of the router unit of Fig. 19A;

Fig. 20A is a block diagram the construction of one of the six input ports of the router unit shown in Fig. 19A;

Fig. 20B is a block diagram of the synchronization logic used to validate command/data symbols received at an input port of the router unit of Fig. 19A;

Fig. 21A is a block diagram illustration of the target port selection... is a block diagram illustration of one of the six output ports of the router unit shown in Fig. 19A;

Fig. 23 is an illustration of the method used to transmit identical information to a duplexed pair CPUs of Fig. 2 in synchronized fashion when the processing system is operating in lock-step (duplex) mode, using a pair the FIFOs of Fig... is a simplified block diagram illustrating the clock generation system of each of the sub-processing systems of Figs. 1A - 1C for developing the plurality of clock signals used to operate the various elements of that sub-processing system;

Fig. 25 illustrates the topology used to interconnect the clock generation systems of paired sub-processing systems for synchronizing the various clock signals of the pair of sub-processing systems to one another;

Fig. 26A and 26B illustrates a FIFO constant rate clock control logic used to control the clock synchronization FIFO of Figs. 8 or 20 in the situation when the two clocks used to... structure of the on-line access port (OLAP) used to provide access to the maintenance

processor (MP) to the various elements of the system of Fig. 1A (or those of Figs... the soft-flag logic used to handle asymmetric variables between the CPUs of paired sub-processing systems operating in duplex mode;

Fig. 31A shows a flow diagram, and Fig. 31B illustrates a portion of SYNC CLK, both of which are used to reset and synchronize the clock synchronization FIFOs of the CPUs and routers of the processing system of Fig. 1A that receive information from each other;

Fig. 32 is a flow... 33A - 33D generally illustrate the procedure used to bring an one of the CPUs of processing system shown in Fig. 1A into lock-step, duplex mode operation with the other of the CPUs without measurably halting operation of the processing system; and

Fig. 34 illustrates a reduced cost architecture incorporating teachings of the invention; and... to the figures and, for the moment, principally Fig. 1A, there is illustrated a data processing system, designated with the reference 10, constructed according to the various teachings of the present invention. As Fig. 1A shows, the data processing system 10 comprises two sub-processor systems 10A and 10B each of which are substantially the same in structure and function... should be appreciated that, unless noted otherwise, a description of any one of the sub-processor systems 10 will apply equally to any other sub-processor system 10.

Continuing with Fig. 1A therefore, each of the sub-processor systems 10A, 10B is illustrated as including a central processing unit (CPU) 12, a router 14, and a plurality of input/output (I/O) packet interfaces... one of the I/O packet interfaces 16 will also have coupled thereto a maintenance processor (MP) 18.

The MP 18 of each sub-processor system 10A, 10B connects to each of the elements of that sub-processor system via an IEEE 1149.1 test bus 17 (shown in phantom in Fig. 1A... accompanying clock signal. As Fig. 1A further

illustrates, TNet Links L also interconnect the sub-processor systems 10A and 10B to one another, providing each sub-processor system 10 with access to the I/O devices of the other as well as inter-CPU communication. As will be seen, any CPU 12 of the processing system 10 can be given access to the memory of any other CPU 12, although... the memory of a CPU 12 by a wayward peripheral device 17.

Preferably, the sub-processor systems 10A/10B are paired as illustrated in Fig. 1A (and Figs 1B and 1C, discussed below), and each sub-processor system 10A/10B pair (i.e., comprising a CPU 12, at least one router 14... 12A) connects, by a TNet Link L to a router (14A) of the corresponding sub-processor system (e.g., 10A). Conversely, the Y port connects the CPU (12A) to the router (14B) of the companion sub-processor system (10B). This latter connection not only provides a communication path for access by a CPU (12A) to the I/O devices of the other sub-processor system (10B), but also to the CPU (12B) of that system for inter-CPU communication.

Information is communicated between any element of the processing system 10 and any other element (e.g., CPU 12A of sub-processor system 10A) of the system and any other element of the system (e.g., an I/O device associated with an I/O packet interface 16B of sub-processor system 10B) via message "packets." Each message packet is made up of a number of... this reason, a unique method of receiving the symbols at the receiver, using a clock synchronization first-in-first-out (CS FIFO) storage structure (described more fully below), has been developed... operation means just that: the frequencies of the clock signals of the transmitter and receiver units are locked, although not necessarily in phase. Frequency locked clock signals are used to transmit symbols between the routers 14A, 14B and the CPUs 12 of paired sub-processor systems (e.g., sub-processor systems 10A, 10B, Fig. 1A). Since the clocks of the transmitting and receiving element are not phase related, a clock synchronization FIFO is again used --albeit operating in a slightly different mode from that used for... difference, as will be seen, is due to the fact that pairs of the sub-processor systems 10 can be operated in a synchronized, lock-step mode, called duplex mode, in which each CPU 12 operates to execute the... 1A illustrates another feature of the invention: a cross-link connection between the two sub-processor systems 10A, 10B through the use of additional routers 14 (identified in Fig. 1A as... RY(sub(1)), and RY(sub(2)) form a cross-link connection between the sub-processors 10A, 10B (or, as shown, "sides" X and Y, respectively) to couple them to L... the routers RX(sub(2)) and RY(sub(2)) provide the I/O packet interface units 16x and 16y with a dual ported interface. Of course, it will now be evident... lend themselves to being used in a manner that can extend the configuration of the processing system 10 to include additional sub-processor systems such as illustrated in Figs. 1B and 1C. In Fig. 1B, for example, one... of each of the routers 14A and 14B is used to connect the corresponding sub-processor systems 10A and 10B to additional sub-processor systems 10A' and 10B' forming thereby a larger processing system comprising clusters of the basic processing system 10 of Fig. 1.

Similarly, in Fig. 1C the above concept is extended to form an eight sub-processor system cluster, comprising sub-processor systems pairs 10A/10B, 10A'/10B', 10A''/10B'', and 10A'''/10B'''. In turn, each of the sub-processor systems (e.g., sub-processor system 10A) will have essentially the same basic minimum configuration of a CPU 12, a... by a I/O packet interface 16, except that, as Fig. 1C shows, the sub-processor systems 10A and 10B include additional routers 14C and 14D, respectively, in order to extend the cluster beyond sub-processor systems 10A'/10B' to the sub-processor systems 10A''/10B'' and 10A'''/10B'''. As Fig. 1C further illustrates, unused ports 4 and... the routers 14 when configuring the topology of the system 10, any CPU 12 of processing system 10 of Fig. 1C can access any other "end unit" (e.g., a CPU or I/O device) of any of the other sub-processor systems. Two paths are available from any CPU 12 to the last router 14 connecting to the I/O packet interface 16. For example, the CPU 12B of the sub-processor system 10B' can access the I/O 16''' of sub-processor system 10A''' via router 14B (of sub-processor system 10B'), router 14D, and router 14B (of sub-system 10B''') and, via link LA... 10A'''), OR via router 14A (of sub-system 10A'), router 14C, and router 14A (sub-processor system 10A'''). Similarly, CPU 12A of sub-processor system 10A'' may access (via two paths) memory contained in the CPU 12B of sub-processor 10B to read or write data. (Memory accesses by one CPU 12 of another component of the processing system requires, as will be seen, the components seeking access to have authorization to do... prevents corruption of memory data of a CPU by erroneous access.)

The topology of the processing system shown in Fig. 1B is achieved by using port 1 of the routers 14A, 14B, and auxiliary TNet links LA, to connect to the routers 14A', 14B' of sub-processor systems 10A', 10B'. The topology thereby obtained establishes redundant communication paths between any CPU 12 (12A, 12B, 12A', 12B') and any I/O packet interface 16 of the processing system 10 shown in Fig. 1B. For example, the CPU 12A' of the sub-processor system 10A' may access the I/O 16A of sub-processor system 10A by a first path formed by the router 14A' (in port 4, out... ..shown in Fig. 1B. By interconnecting one port of each router 14 of each sub-processor pair, and using additional auxiliary TNet links LA (illustrated in Fig. 1C with the dotted line connections) between the ports 1 of the routers 14 (14A'' and 14B'') of sub-processor systems 10A'', 10B'' and 10A''', 10B''', two separate, independent data paths can be found between any CPU 12 and any I/O packet interface 16. In this fashion, any end unit (i.e., a CPU 12 or an I/O packet interface 16) will have at least two paths to any other end unit.

Providing alternate paths of access between any two end units (e.g., between a CPU 12 and any other CPU 12, or between any CPU... ..any two of the remaining fault domains. Here, a fault domain could be a sub-processor system (e.g., 10A). Thus, if the sub-processor system 10A were brought down because of a failure the electrical power being supplied, without... ..TNet link LA between the routers 14A''' and 14B''', the CPU 12B of the sub-processor system 10B would have lost access to the I/O packet interface 16''' (via router... ..with the loss of the router 14A (and router 14C) by loss of the sub-processor system 10A, communications between the CPU 12B is still possible via the route of router... ..equally to CPU 12B. As Fig. 2 shows, the CPU 12A includes a pair of processor units 20a, 20b that are configured for synchronized, lock-step operation in that both processor units 20a, 20b receive and execute identical instructions, and issue identical data and command outputs, at substantially the same moments in time. Each of the processor units 20a and 20b is connected, by a bus 21 (21a, 21b) to a corresponding cache memory 22. The particular type of processor units used could contain sufficient internal cache memory so that the cache memory 22 would not... ..22 could be used to supplement any cache memory that may be internal to the processor units 20. In any event, if the cache memory 22 is used, the bus 21 is... ..22 address bits, 3 bits of parity covering the address, and 7 control bits.

The processors 20a, 20b are also respectively coupled, via a separate 64-bit address/data bus 23 to X and Y interface units 24a, 24b. If desired, the address/data communicated on each bus 23a, 23b could also be protected by parity, although this will increase the width of the bus. (Preferably, the processors 20 are constructed to include RISC R4000 type microprocessors, such as are available from the MIPS Division of Silicon Graphics, Inc. of Santa Clara, California.)

The X and Y interface units 24a, 24b operate to communicate data and command signals between the processor units 20a, 20b and a memory system of the CPU 12A, comprising a memory controller (MC... ..MC halves 26a and 26b) and a dynamic random access memory array 28. The interface units 24 interconnect to each other and to the Mcs 26a, 26b by a 72-bit... ..accompanied by 8 bits of ECC) are written to the memory 28 by the interface units 24, one interface unit 24 will drive only one word (e.g., the 32 most significant portion) of the doubleword being written while the other interface unit 24 writes the other word of the double word (e.g., the least significant 32-bit portion of the doubleword). In addition, on each write operation the interface units 24a, 24b perform a cross-check operation on the data not written by that interface unit 24 with the data written by the other to check for errors; on read operations... ..accessed corresponds to the address of the location from which the doubleword was stored.

Interface units 24a, 24b of the CPU 12A form the circuitry to respectively service the X and Y (I/O) ports of the CPU 12A. Thus, the X interface unit 24a connects by the bi-directional TNet Link Lx to a port of the router 14A of the processor system 10A (Fig. 1A) while the Y interface unit 24b similarly connects to the router 14B of the processor system 10B by TNet Link Ly. The X interface unit 24a handles all I/O traffic between the router 14A and the CPU 12A of the sub-processor system 10A. Likewise, the Y interface unit 24b is responsible for all I/O traffic between the CPU 12A and the router 14B of companion sub-processor system 10B.

The TNet Link Lx connecting the X interface unit 24a to the router 14A (Fig. 1) comprises, as above indicated, two 10-bit buses... sub(x)) carries data incoming from the router 14A. In similar fashion, the Y interface unit 24b is connected to the router 14B (of the sub-processor system 10B) by two 10-bit buses: 30(sub(y)) (for outgoing transmissions) and 32... y)) (for incoming transmissions), together forming the TNet Link Ly.

The X and Y interface units 24a, 24b are synchronously operated in lock-step, performing substantially the same operations at substantially the same times. Thus, although only the X interface unit 24a actually transmits data onto the bus 30(sub(x)), the same output data is being produced by the Y interface unit 24b, and used for error-checking. The Y interface unit 24b output data is coupled to the X interface unit 24a by a cross-link 34(sub(y)) where it is received by the X interface unit 24a and compared against the same output data produced by the X interface unit. In this way the outgoing data made available at the X port of the CPU... the port of the CPU 12A is checked. The output data from the Y interface unit 24b is coupled to the Y port by a 10-bit bus 30(sub(y)), and also to the X interface unit 24a by the 9-bit cross-link 34(sub(y)) where it is checked with that produced by the X interface unit.

As mentioned, the two interface units 24a, 24b operate in synchronous, lock-step with one another, each performing substantially the same... X and/or Y ports of the CPU 12A must be received by both interface units 24a, 24b to maintain the two interface units in this lock-step mode. Thus, data received by one interface unit 24a, 24b is passed to the other, as indicated by the dotted lines and 9... sub(x)) (communicating incoming data being received at the X port by the X interface unit 24a to the Y interface unit 24b) and 36(sub(y)) (communicating data received at the Y port by the Y interface unit 24b to the X interface unit 24a).

Certain more robust operating systems are structured with a fault-tolerant capability in the... example, U.S. Patent No. 4,817,091 teaches a multiprocessor system in which each processor periodically messages each of the processors of the system (including itself), under software control, to thereby provide an indication of continuing operation. Each of the processors, in addition to performing its normal tasks, operates as a backup processor to another of the processors. In the event one of the backup processors fails to receive the messaged indication from a sibling processor, it will take over the operation of that sibling (now thought to be inoperative), in... platform for both types of software. Thus, when a robust operating system is available, the processing system 10 can be configured to operate in a "simplex" mode in which each of... left, in most instances, to software.

Alternatively, for less robust operating systems and software, the processing system 10 provides a hardware-based fault-tolerance by being configured to operate in a... g., CPUs 12A, 12B) are coupled together as shown in Fig. 1A, to operate in synchronized, lock-step fashion, executing the same instructions at the substantially the same moment in time... data and command symbols. In order to simplify the design of the CPU 12, the processors 20 are precluded from communicating directly with any outside entity (e.g., another CPU 12... O device via the I/O packet interface 16). Rather, as will be seen, the processor will construct a data structure in memory and turn over control to the interface units 24. Each interface unit 24 includes a block transfer engine (BTE; Fig. 5) configured to provide a form of... to the destination according to information contained in the message packet.

The design of the processing system 10 permits a memory 28 of a CPU to be read or written by... via the routers 14. Accordingly, before continuing with the description of the construction of the processing system 10, it would be of advantage to understand first the configuration of the data... information.

As indicated, the HADC message packet operates to communicate write data between the end units (e.g., CPU 12) of the processing system 10. Other message packets, however, may be differently constructed because of their function and... CRC. The HC message packet is used to acknowledge a request to write data.

Interface Unit:

The X and Y interface units 24 (i.e., 24a and 24b - Fig. 2) operate to perform three major functions within the CPU 12: to interface the processors 20 to the memory 28; to provide an I/O service that operates transparently to, but under the control of, the processors; and to validate requests for access to the memory 28 from outside sources.

Regarding first the interface function, the X and Y interface units 24a, 24b operate to respectively communicate processors 20a, 20b to the memory controllers (Mcs 26a, 26b) and memory 28 for writing and... ..fast checking of the data read/written. For example, write operations have the two interface units 24a, 24b cooperating to cross-check the data to be written to ensure its integrity (and at the same time, the interface units 24 will operate) to develop an error correcting code (ECC) that covers, as will be... ..to have been retrieved from the appropriate address.

With respect to I/O access, the processors 20 are not provided with the ability to communicate directly with the input/output systems... ..must write data structures to the memory 28 and then pass control to the interface units 24 which perform a direct memory access (DMA) operation to retrieve those data structures, and... ..indicated in the data structure itself.)

The third function of the X and Y interface units 24, access validation to the memory 28, uses an address validation and translation (AVT) table maintained by the interface units. The AVT table contains an address for each system component (e.g., an I/O... ..the incoming message packets are virtual addresses. These virtual addresses are translated by the interface unit to physical addresses recognizable by the memory control units 26 for accessing the memory 28.

Referring to Fig. 5, illustrated is a simplified block diagram of the X interface unit 24a of the CPU 12A. The companion Y interface unit 24b (as well as the interface units 24 of the CPU 12B, or any other CPU 12) is of substantially identical construction. Accordingly, it will be understood that a description of the interface unit 24a will apply equally to the other interface units 24 of the processing system 10.

As Fig. 5 illustrates, the X interface unit 24a includes a processor interface 60, a memory interface 70, interrupt logic 86, a block transfer engine (BTE) 88, access validation and translation logic 90, a packet transmitter 94, and a packet receiver 96.

Processor Interface:

The processor interface 60 handles the information flow (data and commands) between the processor 20a and the X interface unit 24a. A processor bus 23, including a 64 bit address and data bus (SysAD) 23a and a 9 bit command bus 23b, couples the processor 20a and the processor interface 60 to one another. While the SysAD bus 23a carries memory address and data... ..and qualifying commands carried at substantially the same time on the SysAD bus 23a. The processor interface 60 operates to interpret commands issued by the processor unit 20a in order to pass reads/writes to memory or control registers of the processor interface. In addition, the processor interface 60 contains temporary storage (not shown) for buffering addresses and data for access to... ..26). Data and command information read from memory is similarly buffered en route to the processor unit 20a, and made available when the processor unit is ready to accept it. Further, the processor interface 60 will operate to generate the necessary interrupt signalling for the X interface unit 24a.

The processor interface 60 is connected to a memory interface 70 and to configuration registers 74 by a bi-directional 64 bit processor address/data bus 76. The configuration registers 74 are a symbolic representation of the various control registers contained in other components of the X interface unit 24a, and will be discussed when those particular components are discussed. However, although not specifically... ..throughout other of the logic that is used to implement the X interface 24a, the

processor address/data bus 76 is likewise coupled to read or write to those registers.

Configuration registers 74 are read/write accessible to the processor 20a; they allow the X interface unit to be "personalized." For example, one register identifies the node address of the CPU 12A... ..with the CPU 12A; another, readable only, contains a fixed identification number of the interface unit 24, and still other registers define areas of memory that can be used by, for... ..logic 90, etc.) employing them are discussed.

The memory interface 70 couples the X interface unit 24a to the memory controllers 26 (and to the Y interface unit 24b; see fig. 2) by a bus 25 that includes two 36 bi-directional bit... 25a, 25b. The memory interface operates to arbitrate between requests for memory access from the processor unit 20, the BTE 88, and the AVT logic 90. In addition to memory accesses from the processor unit 20a, the memory 28 may also be accessed by components of the processing system 10 to, for example, store data requested to be read by the processor unit 20a from an I/O unit 17, or memory 28 may also be accessed for I/O data structures previously set up in memory by the processor unit. Since these accesses are all asynchronous, they must be arbitrated, and the memory interface 70... command information accessed from the memory 28 is coupled from the memory interface to the processor interface 60 by a memory read bus 82, as well as to an interrupt logic... doubleword quantities. However, while the memory interfaces 70 of both the X and Y interface units 24a and 24b formulate and apply the (64-bit) doubleword to the bus 25, each... by the memory interface 70 are coupled to the memory interface by the companion interface unit 24 where they are compared with the same 32 bits for error.

Digressing for the... containing interrupt information are received, that information is conveyed to the interrupt logic 86 for processing and posting for action by the processor 20, along with any interrupts generated internal to the CPU 12A. Internally generated interrupts will... register 71 (internal to the interrupt logic 86), indicating the cause of the interrupt. The processor 20 can then read and act upon the interrupt. The interrupt logic is discussed more fully below.

The BTE 88 of the X interface unit 24a operates to perform direct memory accesses, and provides the mechanism that allows the processors 20 to access external resources. The BTE 88 can be set-up by the processors 20 to generate I/O requests, transparent to the processors 20 and notify the processors when the requests are complete. The BTE logic 88 is discussed further below.

Requests for... 8 byte wide format necessary for storing in the memory 28.

Outgoing message packets containing processor originated transaction requests (e.g., a read request asking for a block data from an I/O unit) are monitored by the request transaction logic (RTL) 100. The RTL 100 provides a time... will generate an interrupt (handled and reported by the interrupt logic 86) to inform the processor 20 that the request was not honored. In addition, the RTL 100 will validate responses... 28 (by the DMA operation of the BTE 86) at a location known to the processor 20 so that it can locate the response.

Each of the CPUs 12 are checked... discussed. One such check is an on-going monitor of the operation of the interface units 24a, 24b of each CPU. Since the interface units 24a, 24b operate in lock-step synchronism checking can be performed by monitoring the operating states of the paired interface units 24a, 24b by a continuous comparison of certain of their internal states. This approach is implemented by using one stage of a state machine (not shown) contained in the unit 24a of CPU 12A, and comparing each state assumed by that stage with its identical state machine stage in the interface unit 24b. All units of the interface units 24 use state machines to control their operations. Preferably, therefore, a state machine of the memory interface 70 that controls the data transfers between the interface unit 24 and the MC 26 is used. Thus, a selected stage of the state machine used in the memory interface 70 of the interface unit 24a is selected. An identical stage of a state machine of one of the interface unit 24b is also selected. The two selected stages are communicated between the interface units 24a, 24b and received by a compare circuit contained in both interface units 24a, 24b. As the interface units operate lock-step with one another, the state machines will likewise march through the same identical states, assuming each state at substantially the same moments in time. If an interface unit encounters an error, or fails, that activity will cause the interface units to diverge, and the state machines will assume different states. The time will come when... that will bring to the attention of the CPUs 12A (or 12B) that the interface units 24a, 24b of that CPU are no longer in lock-step, and to act accordingly... X port, receiving only those message packets transmitted by the router 14A of the sub-processor system 10A (Fig. 1A). The Y port is serviced by the Y interface unit 24b to receive message packets from the router 14B of the companion sub-processor system 10B. However, both interfaces (as well as Mcs 26 and processor 20), as has been indicated, are basically mirror images of one another in that both... in both structure and function. For

this reason, message packet information, received by one interface unit (e.g., 24a) must be passed for processing also to the companion interface unit (e.g., 24b). Further, since both interface units 24a, 24b will assemble the same message packets for transmission from the X or the Y ports, the message packet being transmitted by the interface unit (e.g., 24b) actually being communicated from the associated port (e.g., the Y port) will also be coupled to the other interface unit (e.g., 24a) for cross-checking for errors. These features are illustrated in Figs. 6... ..receiving portions of the packet receivers 96 (96x, 96y) of the X and Y interface units 24a, 24b are broadly illustrated. As shown, each packet receiver 96x, 96y has a clock... ..receive a corresponding one of the TNet Links 32. The CS FIFOs 102 operate to synchronize the incoming command/data symbols to the local clock of the packet receiver 96, buffering... ..104x, coupled to the MUX 104y of the packet receiver 96y of the Y interface unit 24b by the cross-link connection 36(sub(x)). In similar fashion, information received at the Y port is coupled to the X interface unit 24a by the cross-link connection 36(sub(y)). In this manner, the command/data... ..packets received at one of the X, Y ports by the corresponding X, Y, interface unit 24a, 24b is passed to the other so that both will process and communicate the same information on to other components of the interface units 24 and/or memory 28.

Continuing with Fig. 6, depending upon which port X, Y... ..or the other of the CS FIFOs 102x, 102y for communication to the storage and processing logic 110 of the interface unit 24. The information contained in each 9-bit symbol is an 8-bit byte of... ..the encoding of which is discussed below with respect to Fig. 9. The storage and processing logic 110 will first translate the 9-bit symbols to 8-bit data or command... ..the outputs of the CS FIFOs 102x, 102y are also coupled to a command decode unit in addition to the MUX 104. The command decode unit operates to recognize command symbols (differentiating them from data symbols in a manner that is... ..below), decoding them to generate therefrom command signals that are applied to a receiver control unit, a state machine-based element that functions to control packet receiver operations.

As indicated above... ..at the output of the MUX 104, the receiver control portion of the storage control unit enables CRC check logic 106 to calculate a CRC symbol while the data symbols are... ..below, CS FIFOs are found not only in the packet receivers 96 of the interface units 24, but also at each receiving port of the routers 14 and the I/O... ..an even more important part, and perform a unique function, when a pair of sub-processor systems are operating in duplex mode and the two CPUs 12A and 12B of the sub-processor systems 10A, 10B operate in synchronized, lock-step, executing the same instructions at the same time. When operating in this latter... ..difficult to ensure that the clocking regime of the routers 14A and 14B are exactly synchronized to those of the CPUs 12A and 12B - even when using frequency locked clocking. In... ..used to transmit symbols to a CPU 12 and the clock used by an interface unit 24 to receive those symbols.

The structure of the CS FIFO 102 is diagrammatically illustrated... ..i.e., a packet) or IDLE symbols - except during certain situations (e.g., reset, initialization, synchronization and others discussed below). As explained above, each symbol held in the transmit register 120... ..same symbol leaving the storage queue, allowing each symbol entering the storage queue 126 to settle before it is clocked out and passed to the storage and processing units 110x (and 110y) by the MUX 104x (and 104y). Since the transmit and receive clock synchronization FIFOs are used at these other ports to receive symbols transmitted with near frequency clocking, and the structure of these clock synchronization FIFOs are substantially the same as that used in frequency locked environments, i.e., that... ..of the storage queue 126 are nine bits wide; in near frequency environments, the clock synchronization FIFOs use symbol locations of the queue 126 that are 10 bits wide, the extra... ..the faster clock source. To handle this clock drift, the two pointers are effectively re-synchronized periodically.

When the CPUs 12 are paired and operating in duplex mode, all four interface

units 24 operate in lock-step to, among other things, transmit the same data and receive... ..simplex mode, each independent of the other, clocking need only be near frequency.

The interface unit 24 receives a SYNC CLK signal that is used in combination with a SYNC command symbol to initialize and synchronize the Rcv register 124 to the transmitting router 14. When using either near frequency or... 102X preferably begin from some known state. Incoming symbols are examined by the storage and processing units 110 of the packet receivers 96. The storage and processing units look for, and act upon as appropriate, command symbols. Pertinent here is that when the... receives a SYNC command symbol it will be decoded and detected by the storage and processing unit 110. Detection of the SYNC command symbol by the storage and processing unit 110 causes assertion of a RESET signal. The RESET signal, under synchronous control of the SYNC CLK signal, is used to reset the input buffers (including the clock synchronization buffers) to predetermined states, and synchronize them to the routers 14.

The synchronization of the CS FIFOs 102 of the interface units 24 those of one or both routers 14A, 14B is discussed more fully below in the section discussing synchronization.

Packet Transmitter:

Each interface unit 24 is assigned to transmit from and receive at only one of the X or Y ports of the CPU 12. When one of the interface units 24 transmits, the other operates to check the data being transmitted. This is an important... shows, in abbreviated form, the packet transmitters 94x, 94y of the X and Y interface units 24a, 24b, respectively. Both packet transmitters are identically constructed, so that discussion of one (packet... logic 152 that receives, from the BTE 88 or AVT 90 of the associated interface unit (here, the X interface unit 24a) the data to be transmitted - in doubleword (64-bit) format. The packet assembly logic... and Y ports: they are either symbols that make up a message packet in the process of being transmitted, or IDLE symbols, or other command symbols used to perform control functions... 154, 156. The output of the multiplexer 154 connects to the X port. (The interface unit 24b connects the output of the multiplexer 154 to the Y port.) The multiplexer 156... sub(x)) to the checker logic 160 of the packet transmitter 94y (of the interface unit 24b).

A selection (S) input of the multiplexers receives a 1-bit output from an... is accessible to the MP 18 via an OLAP (not shown) formed in the interface unit 24, and is written with information that "personalizes," among other things, the interface units 24. Here, the X/Y stage of the configuration register 162 configures the packet transmitter 94x of the X interface unit 24a to communicate the X encoder 150x output to the X port; the output of... traffic is present, the operation of the two packet interfaces 94 (and, thereby, the interface units 24 with which they are associated) are continually monitored. Should one of the checkers detect... will be asserted, resulting in an internal interrupt being posted for appropriate action by the processors 20.

Message packet traffic operates in the same manner. Assume, for the moment, that the... that information, a byte at a time, to the X encoder 150x of both interface units 96, which will translate each byte to encoded 9-bit form. The output of the... is checked with that from the packet transmitter 94x. Again, the operation of the interface units 24a, 24b, and the packet transmitters they contain, are inspected for error.

In the same... monitored.

Returning for the moment to Fig. 5, if the outgoing message packet is a processor initiated transaction (e.g., a read request), the processors 20 will expect a message packet to be returned in response. Thus, when the BTE... will issue a timeout signal to the interrupt logic (Fig. 14A) to thereby notify the processors 20 of the absence of a response to a particular transaction (e.g., a read... the access, to name just a few. Also, the area of memory of the memory unit 28 desired to be accessed are identified in the message packets by virtual or I... virtual addresses be translated to physical addresses of the memory 28. Finally, interrupts generated by units or elements external to the CPU 12A, are transmitted via message packets to interrupt the processors 20, which are also written to memory 28 when received. All this is handled by the interrupt logic and AVT logic 86, 90.

The AVT logic unit 90 utilizes a table (maintained by the processor 20 in memory 28) containing AVT entries for each possible external source permitted access to the memory 28. Each AVT entry identifies a specific source

element or unit and the particular page (a page being nominally 4K (4096) bytes), or portion of a... ..expected" memory accesses. Expected memory accesses are those initiated by the CPU 12 (i.e., processors 20) such as a read request for information from an I/O device. These latter memory accesses are handled by a transaction sequence number (TSN) assigned to each processor initiated request. At about the time the read request is generated, the processors 20 will allocate an area of memory for the data expected to be received in... ..and 26b are, in turn, respectively coupled to the memory interfaces 70 of each interface unit 24a, 24b. The 64-bit doublewords are written to the memory 28 with the upper... ..check bits respectively from the memory interfaces 70 (70a, 70b) of each of the interface units 24a, 24b (Fig. 5).

Referring to Fig. 10, each memory interface 70 receives, from either the bus 82 from the processor interface 60 or the bus 83 from AVT logic 90 (see Fig. 5), of the associated interface unit 24, 64 bits of data to be written to memory. The busses 76 and 83... ..other for cross-checking between them. Thus, for example, the memory interface 70a (of interface unit 24a) will drive the MC 26a with the "upper" 32 bits of the 64 bits... ..are check bits, leaving 40 bits unused.

Access Validation:

As previously indicated, components of the processing system 10 external to the CPU 12A (e.g., devices of the I/O packet... ..not without qualification. Access validation, as implemented by the AVT logic 90 of the interface units 24, operates to prevent the content of the memory 28 from being corrupted by erroneously... ..Accesses to the memory 28 are validated by the AVT logic 90 of each interface unit 24 (Fig. 5), using all of six checks: (1) that the CRC of the message... ..also are permitted the particular message packet source.

The access validation mechanism of the interface unit 24a, AVT logic 88, is shown in greater detail in Fig. 11. Incoming message packets... ..and post an interrupt to the interrupt logic 86 (Fig. 5) for action by the processor 20.

The mask operation permits the size of the table of AVT entries to be varied. The content of the AVT mask register 175 is accessible to the processor 20, permitting the processors 20 to optionally select the size of the AVT entry table. A maximum AVT table... ..172 allows the AVT size to be matched to the needs of the system. A processing system 10 that includes a larger number of external elements (e.g., the number of... ..amount of the memory space of memory 28 to the AVT entries. Conversely, a smaller processing system 10, with a smaller number of external elements will not have such a large... ..set to a logic "ZERO" indicate an nonexistent TNet address, outside the limits of the processing system 10. A received packet with a TNet address outside the allowable TNet range will... ..in Fig. 11 as being held in the AVT entry register 180 during the validation process. AVT entries have two basic formats: normal and interrupt. The format of a normal AVT... ..of the AVT input register 170) will result in an error being posted to the processor via an interrupt.

A 12-bit "Permissions" field is included in the AVT entry to...path=0). Denials are logged as interrupts with the interrupt logic, and reported to the processor 20 - if the E field is set to a state ("ONE") that enables error-reporting... ..e.g., to a "ONE"), the other fields (Upper Bound, etc.) gain new definitions for processing interrupt writes and managing interrupt queues. This is discussed in more detail below in connection... ..memory 28 will be handled. Set to one state, the requested write operation will be processed normally; set to a second state, write requests specifying addresses with a fractional cache line... ..be written to a specific queue (interrupt queue) in memory 28, with signalling provided the processors 20 to indicate that an interrupt has been received and "posted," and ready for servicing by the processors 20. Since the interrupt queues are at specific memory locations, the processor can obtain the interrupt data when needed.

An AVT interrupt entry for an interrupt may... ..by the interrupt logic 86, and extracted from the head of the queue by the processor 20 when servicing the interrupt.

The AVT interrupt entry also includes a 20-bit segment ("Source ID") containing source ID information, identifying the external unit seeking attention by the interrupt process. If the source ID information of the AVT interrupt entry

does not match that contained... ..class" of the interrupt that is used to determine the interrupt level set in the processor 20 (described more fully below); (2) a queue number that is used to select, as... ..capability to deliver interrupts to a CPU 12 for servicing. For example, an I/O unit may be unable to complete a read or write transaction issued by a CPU because... ..identify the recipient. These and other errors, exceptions, and irregularities, noted by the I/O units, or the I/O Interface elements, can become the a condition that requires the intervention... ..the AVT entry register 180 for use by the interrupt logic 86 of the interface unit 24 (Fig. 5), illustrated in greater detail in Fig. 14A.

It is interrupt logic 86... ..four circular queues specified by the base address information contained in the AVT entry. The processor(s) 20 will then be notified, and it will be up to them as to... ..selected tail queue register 256 by combiner circuit 270, the output of which is the processed by the "mod z" circuit 273 to turn new offset into the queue at which... ..signal. The Queue Full warning signal becomes an "intrinsic" interrupt that is conveyed to the processor units 20 as a warning that if the matter is not promptly handled, later-received interrupt... ..will be discarded.

Incoming message packet interrupts will cause interrupts to be posted to the processor 20 by first setting one of a number of bit positions of an interrupt register 280. Multi-entry queued interrupts are set in interrupt registers 280a for posting to the processor 20; single-entry queue interrupts use interrupt register 280b. Which bit is set depends upon... ..multi-entry queued interrupts, soon after a multi-entry queued interrupt is determined, the interface unit will assert a corresponding interrupt signal (I1) that is applied to decode circuit 283. Decode... ..of register 280a to set, thereby providing advance information concerning the received interrupt to the processor(s) 20, i.e., (1) the type of interrupt posted, and (2) the class of... ..to one another by a compare circuit 279. The update register is writable by the processor 20 to select a register pair for comparison. If the content of the two selected... ..cleared.

Digressing for the moment, there are two basic types of interrupts that concern the processors 20: those interrupts that are communicated to the CPU 12 by message packets, and those... ..the seven interrupt postings to a latch 288, from which they are coupled to the processor 20 (20a,20b) which has an interrupt register for receiving holding the postings.

In addition... ..change in interrupts (either an interrupt has been serviced, and its posting deleted by the processor 20, or a new interrupt has been posted), a "CHANGE" signal will be issued to the processor interface 60 to inform it that an interrupt posting change has occurred, and that it should communicate the change to the processor 20.

Preferably, the AVT entry register 180 is configured to operate like a single line... ..such as set-associative, fully-associate, or direct-mapped, to name a few.

Coherency:

Data processing systems that use cache memory have long recognized the problem of coherency: making sure that... ..the incoming packet is permitted access are applied to a boundary crossing (Bdry Xing) check unit 219. Boundary check unit 219 also receives an indication of the size of the cache block the CPU 12... ..Len field of the header information from the AVT input register 170. The Bdry Xing unit determines if the data of the incoming packet is not aligned on a cache boundary... ..time an interrupt will be written to the queued interrupt register 280, to alert the processors 20 that a portion of the incoming data is located in the special queue.

In... ..not, the packet (both header and data) is written to a special queue, and the processors so notified by the intrinsic interrupt process described above. The processors may then move the data from the special queue to cache 22, and later write... ..the cache 22 and the memory 28 is preserved.

Block Transfer Engine (BTE):

Since the processor 20 is inhibited from directly communicating (i.e., sending) information to elements external to the... ..indirect method of information transmission.

The BTE 88 is the mechanism used to implement all processor initiated I/O traffic to transfer blocks of information. The BTE 88 allows creation of... ..BTE registers 300, 302 whose content is coupled to the MUX 306 (of the interface unit 24a; Fig. 5) and used to access the system memory 28 via the memory controllers... ..BTE data structure 304 in the memory 28 of the CPU 12A (Fig. 2). The processors 20 will write a data structure 304 to the memory 28 each time information is... ..begin on a quadword boundary, and the BTE registers 300, 302 are writable by the processors 20 only. When a processor does write one of the BTE registers 300, 302, it does so with a word... ..the request bit (rc0, rc1) to a clear state, which operates to initiate the BTE process, which is controlled by the BTE state machine 307.

The BTE registers 300, 302 also... ..cause (ec) bit differentiates time-outs and NAKs.

When information is being transferred by the processors 20 to an external unit, the data buffer...the data structure 304 holds the information to be transferred. When information from an external unit is received by the processors 20, the data buffer portion 304b is the location targeted to hold the read response information.

The beginning of the data structure 304, portion 304a written by the processor 20, includes an information field (Dest), identifying the external element which will receive the packet... ..the transmitted data is to be written. This information is used by the packet transmitter unit 120 (Fig. 5) to assemble the packet in the form shown in Figs. 3-4... ..list (el) bit, when set, indicates the end of the chain, and halts the BTE processing.

The interrupt completion (ic) bit, when set, will cause the interface unit 24a to assert an interrupt (BTECmp) which sets a bit in the interrupt register 280... ..the chain pointer).

The interrupt time-out (it) bit, when set, will cause the interface unit 24a to assert an interrupt signal for the processor 20 if the acknowledgement of the access times-out (i.e., if the request timer... ..time), or elicits a NAK response (indicating that the target of the request could not process the request).

Finally, if the check sum (cs) bit is set, the data to be... ..containing the data from which the check sum was formed.

To sum up, when the processors 20 of the CPU 12A desire to send data to an external unit, they will write a data structure 304 to the memory 28, comprising identifier information in portion 304a of the data structure, and the data in the buffer portion 304b. The processors 20 will then determine the priority of the data and will write the BTE register... ..information, and sent.

If the data structure 304 indicates a read request (i.e., the processors 20 are seeking data from an external unit - either an I/O device or a CPU 12), the Len and Local Buffer Ptr... ..receiver 100 (Fig. 5) until the local memory write operation is executed.

Responses to a processor-generated read request to an external unit are not processed by the AVT table logic 146. Rather, when the processors 20 set up the BTE data structure, a transaction sequence number (TSN) is assigned the... ..the BTE 88, which will be an HAC type packet (Fig. 4) discussed above. The processors 20 will also include an memory address in the BTE data structure at which the... ..302, assume that the foregoing transfer of data from the CPU 12A to an external unit is of a large block of information. Accordingly, a number of data structures would be set up in memory 28 by the processors 20, each (except the last) including a chain pointer to additional data structures, the sum... ..sent. Assume now that a higher priority request is desired to be made by the processors 20. In such a case, the associated data structure 304 for such higher priority request... ..with another BTE operation descriptor.

Memory Controller:

Returning, for the moment, to Fig. 2, interface units 24a, 24b access the memory 28 via a pair of memory controllers (MC) 26a, 26b. The Mcs provide a fail-fast interface between the interface units 24 and the memory 28. The Mcs 26 provide the control logic necessary for accessing... ..in dynamic random access memory (DRAM) logic). The Mcs receive memory requests from the interface units 24, and execute reads and writes as well as providing refresh signals to the DRAMs... ..to provide a 72 bit data path between the memory array 28 and the interface units 24a, 24b, which utilize an SBC-DBD-SbD ECC scheme, where $b=4$, on a... ..26a, 26b to work together and simultaneously supply a 64-bit word to the interface units 24 with minimum latency, one-half of which (DO) comes from the MC 26a, and the other half (D1) comes from the other MC 26b. The interface unit 24 generate and check the ECC check bits. The ECC scheme used will not only... ..26 bus 25, as well as in internal registers.

From the viewpoint of the interface units 24, the memory 28 is accessed with two instructions: a "read N doubleword" and a... ..doubleword read or a block read format. The signal called "data valid" tells the interface units 24 two cycles ahead of time that read data is being returned or not being returned.

As indicated above, the maintenance processor (MP 18; Fig. 1A) has two means of access to the CPUs 12. One is... ..18 will write a register contained in the OLAP 285 with instructions that permit the processors 20 to build an image of a sequence of instructions in the memory that will permit them (the processors 20) to commence operation, going to I/O for example to transfer instructions and data from an external (storage) device that will complete the boot process.

The OLAP 285 is also used by the processors 20 to communicate to the MP 18 error indications. For example, if one of the interface units 24 detect a parity error in data received from the memory controller 26, it will... ..and address transfers on the bus 25 between the MC 26a and the corresponding interface unit 24a. The addressing and data transfers on the DRAM data bus, as well as generation... ..the CPU 12.

Packet Routing:

The message packets communicated between the various elements of the processing system 10 (e.g., CPUs 12A, 12B, and devices coupled to the I/O packet... ..First, each TNet Link L connects to an element (e.g., router 14A) of the processing system 10 via a port that has both receive and transmit capability. Each transmit port... ..cycle (i.e., each clock period) of the T(underscore)Clk so that the clock

synchronization FIFO at the receiving end of the transmission will maintain synchronization.

Clock synchronization is dependent upon the mode in which the processing system 10 is operated. If operating in the simplex mode in which the CPUs 12A... ..connect directly to the CPUs may drift with respect to each other. Conversely, when the processing system 10 operates in a duplex mode (e.g., the CPUs operate in synchronized, lock-step operation), the clocks between routers 14 and the CPUs 12 to which they... ..not necessarily phase-locked).

The flow of data packets between the various elements of the processing system 10 is controlled by command symbols, which may appear at any time, even within... ..initiated by a CPU 12, or MP 18, and promulgated to all elements of the processing system 10 by the routers 14 to communicate an event requiring software action by all... ..command symbol is used in conjunction with near frequency operation as an aid to maintaining synchronization between the two clock signals that (1) transfer each symbol to, and load it in each receiving clock synchronization FIFO, and (2) that retrieves symbols from the FIFO.

SLEEP: This command symbol is sent by any element of the processing system 10 to indicate that no additional packet (after the one currently being transmitted, if... ..received.

SOFT RESET (SRST): The SRST command symbol is used as a trigger during the processes ("synchronization" and "reintegration," described below) that are used to synchronize symbol transfers between the CPUs 12 and the routers 14A, 14B, and then to place... ..SYNC command symbol is sent by a router 14 to the CPU 12 of the

processing system 10 (i.e., the sub-processor systems 10A/10B) to establish frequency-lock synchronization between CPUs 12 and routers 14A, 14B prior to entering duplex mode, or when in duplex mode to request synchronization, as will be discussed more fully below. The SYNC command symbol is used in conjunction... or duplex to simplex), among other things, as discussed further below in the section on Synchronization and Reintegration.

THIS LINK BAD (TLB): When any system element receiving a symbol from a TNet link L (e.g., a router, a CPU, or an I/O unit) notes an error when receiving a command symbol or packet, it will send a TLB... identical pairs of symbols that are compared to one another when pulled from the clock synchronization FIFOs. The DVRG command symbol signals the CPU 12 that a mis-compare has been noted. When received by the CPUs, a divergence detection process is entered whereby a determination is made by the CPUs which CPU may be failing... command symbols described above operate to control message flow between the various elements of the processing system 10 (e.g., CPUs 12, router 14, and the like), using principally the BUSY... particular TNet port however, an "end node" (i.e., a CPU 12 or I/O unit 17 - Fig. 1) may not assert backpressure because one of its transmit ports is backpressured... Improperly addressed packets are discarded by the router 14.

When a system element of the processing system 10 receives a BUSY command symbol on a TNet link L on which it... other command symbols (READY, BUSY, etc.).

Whenever a TNet port of an element of the processing system 10 detects receipt of a READY command symbol, it will terminate transmission of FILL... receives.

As will be seen, all elements (e.g., router 14, CPUs 12) of the processing system 10 that connect to a TNet link L for receiving transmitted symbols will receive those symbols via a clock synchronization (CS) FIFO. For example, as discussed above, the interface units 24 of CPUs 12 include all CS FIFOs 102x, 102y (illustrated in Fig. 6). The... depth to allow for speed matching, and the elastic FIFOs must provide sufficient depth for processing delays that may occur between transmission of a BUSY command symbol during receipt of a... another data byte in packet B. As packet A progresses to the next router, the process would be repeated. If the router 14 displaces more data bytes than the FIFO can... irrespective of its own findings.

SLEEP Protocol:

The SLEEP protocol is initiated by a maintenance processor via a maintenance interface (an on-line access port - OLAP), described below. The SLEEP protocol process) in order to change modes without causing data loss or corruption. When a SLEEP command symbol is received, the receiving element of processing system 10 inhibits initiation of transmission of any new packet on the associated transmit port... The HALT command symbol provides a mechanism for quickly informing all CPUs 12 in a processing system 10 that is necessary to terminate I/O activity (i.e., message transmissions between... CPUs that receive HALT command symbols on either of their receive ports (of the interface units 24) will post an interrupt to the interrupt register 280 if the system halt interrupt... interrupt; Fig. 14A).

The CPUs 12 may be provided with the ability to disable HALT processing. Thus, for example, the configuration registers 75 of the interface units 24 can include a "halt enable register" that, when set to a predetermined state (eg., ZERO) disables HALT processing, but reporting detection of a HALT symbol as an error.

Router Architecture:

Referring now to... simplified block diagram of the router 14A is illustrated. The other routers 14 of the processing system 10 (e.g., routers 14B, 14', etc.) are of substantially identical construction and, therefore... these ports 4, 5 are structured to operate in a frequency locked environment when a processing system 10 is set for duplex mode operation. In addition, when in duplex mode, a... 5)) will receive the command/data symbols from the CPUs, pass them through the clock synchronization FIFOs 518 (discussed further below), and compare each symbol exiting the clock synchronization FIFOs with a gated compare circuit 517. When duplex operation is entered, a configuration

register... 517 to activate the symbol by symbol comparison of the symbols emanating from the two synchronization FIFOs 518 of the router input logic 502 for the ports 4 and 5. Of... to that received, at substantially the same time, by the other port input.

To maintain synchronization in the duplex mode, the two port outputs of the router 14A that transmit to... mode, are duplicated by the routers 14, and returned to both CPUs.) The output logic units 504(sub(4)), 504(sub(5)) that are coupled directly to the CPUs 12 will... message packet identifies only one of the duplexed CPUs 12, e.g., CPU 12A) in synchronized fashion, presenting those symbols in substantially simultaneous fashion to the two CPUs 12. Of course, the CPUs 12 (more accurately, the associated interface units 24) receive the transmitted symbols with synchronizing FIFOs of substantially the same structure as that illustrated in Fig. 7A so that, even... from the FIFO structures by both CPUs 12 on the same instruction cycle, maintaining the synchronized, lock-step operation of the CPUs 12 required by the duplex operating mode.

As will... conjunction with configuration data written to registers contained in control logic 509 by the maintenance processor 18 (via the on-line access port 285' and serial bus 19A; see Fig. 1A... links L. The input logic 505 of each port input 502 also assists in maintaining synchronization - at least for those ports sending symbols in the near-frequency environment - by removing received... slower-receiving element receiving symbols from a faster-sending element could overload the input clock synchronization FIFO of the slower-receiving element. That is, if a slower clock is used to pull symbols from the clock synchronization FIFO put there by a faster clock, ultimately the clock synchronization FIFO will overflow.

The preferred technique employed here is to periodically insert SKIP symbols in... stream to avoid, or at least minimize, the possibility of an overflow of the clock synchronization FIFO (i.e., clock synchronization FIFO 518; Fig. 20A) of a router 14 (or CPU 12) due to a T... being slightly higher in frequency than the local clock used to pull symbols from the synchronization FIFO. Using SKIP symbols to by-pass a push (onto the FIFO) operation has the... stall each time a SKIP command symbol is received so that, insofar as the clock synchronization FIFO is concerned, the transmitting clock that accompanied the SKIP symbol was missing.

Thus, logic... the port inputs 502 will recognize, and key off receipt of, SKIP command symbols for synchronization in the near frequency clocking environment so that nothing is pushed onto the FIFO, but... 14, or between routers 14, or between a router 14 and an I/O interface unit 16A - Fig. 1) at a 50 Mhz rate, this allows for a worst case frequency... symbol by supplying FILL or IDLE symbols (which are received and pushed onto the clock synchronization FIFOs, but are not passed to the elastic FIFOs). In short, each elastic FIFO 506... received symbols are then communicated from the input register 516 and applied to a clock synchronization FIFO 518, also by the T(underscore)Clk. The clock synchronization FIFO 518 is logically the same as that illustrated in Figs. 8A and 8B, used in the interface units 24 of the CPUs 12. Here, as Fig. 20A shows, the clock synchronization FIFO 518 comprises a plurality of registers 520 that receive, in parallel, the output of... 516. Associated with each of the registers 520 is a two-stage validity (V) bit synchronizer 522, shown in greater detail in Fig. 20B, and discussed below. The content of each... registers 520, together with the one-bit content of each associated two-stage validity bit synchronizer 522, are applied to a multiplexer 524, and the selected register/synchronizer pulled from the FIFO, and coupled to the elastic FIFO 506 by a pair of... is determined the state of the Push Select signal provided by a push pointer logic

unit 530; and, selection of which register 520 will supply its content, via the MUX 524... and loading of the register 520 selected by the push pointer logic 530. Similarly, the synchronization FIFO control logic 534 receives the clock signal local to the router (Rcv Clk) to... pointer logic 532.

Digressing for a moment, and referring to Fig. 20B, the validity bit synchronizer 522 is shown in greater detail as including a D-type flip-flop 541 with... 530 (Fig. 20A) selects the register 520 of the FIFO with which the validity bit synchronizer is associated for receipt of the next symbol - if not a SKIP symbol.

The delay... Truth Table, below). The D-type flip-flop 543 acts as an additional stage of synchronization, ensuring a stable level at the V output relative to the local Rec Clk. The... flip-flop 542, allowing the Pull signal (a periodic pulse from the sync FIFO Control unit 534) to clear the validity bit on this validity synchronizer 522 when the associated register 520 has been read. (Table omitted)

In summary, the validity synchronizer 522 operates to assert a "valid" (V) signal when a symbol is loaded in a... blocked from being routed out a particular port because another message is already in the process of being routed out that port. However, that other message in turn is also blocked... an incoming message packet bound for the CPUs will be replicated by the crossbar logic unit by routing the message packet to both port output 504(sub(4)) and 504(sub... P) identifies which of path (X or Y) should be used for accessing two sub-processing the device.

The routers 14 provide a capability of constructing a large, versatile routing network for, for example, massively parallel processing architectures. Routers are configured according to their location (i.e., level) in the network by...)) and 509(sub(k)) are such that bits "def" are used in the algorithmic process, then bits "abc" of the Region ID are compared to the content of the Device... the route to default register 509(sub(f)) to the final stage of the selection process: check logic 602. Check logic 602 operates to check the status of the port output... a lower level router, and may be located in one or another of the sub-processing systems 10A, 10B. Whether a router is an upper level or lower level router depends... of CPUs 12 and I/O devices 16 to one another, forming a massively parallel processing (MPP) system. Other such MPP systems may exist, and it is those routers configured as... captured. As soon as the message packet's Destination ID is so captured, the selection process begins, proceeding to the development of a target port address that will be used to... an error that will be posted to the MP18 via the router's (or interface unit's) OLAP for action.

Digressing, it should be appreciated that these protocol rules observed by the routers 14 are also observed by the CPUs 12 (i.e., interface units 24) and I/O packet interfaces 17.

Finally, when the router 14A is in the... directly with the CPUs 12A, 12B, and duplex mode is used, a duplex operation logic unit 638 is utilized to coordinate the port output connected to one of the CPUs 12A... was able to write instructions to the OLAP 285 that would be executed by the processors 20 to build a small memory image and routine to permit the CPU 12 to... the clock generation circuit design. There will be one clock generator circuit in each sub-processor system 10A/10B (Fig. 1) to maintain synchronism. Designated generally with the reference numeral 650... used by the various elements (e.g. CPU. 12, routers 14, etc.) of the sub-processor system containing the clock circuit 650 (e.g., 10A).

The clock generator 654 is shown... The 50 Mhz clock signals produced by the counter 663 are distributed throughout the sub-processor system where needed.

Turning now to Fig. 25, there is illustrated the interconnection and use... the clock circuits 650 used to develop synchronous clock signals for a pair of sub-processor systems 10A, 10B (Fig. 1) for frequency locked operation. As illustrated in Fig. 25, the two CPUs 12A and 12B of the sub-processor systems 10A, 10B each have a clock circuit 650, shown in Fig. 25 as clock... 654B of both CPUs 12. A driver and signal line 667 interconnects the two sub-processor systems to deliver the M(underscore)CLK signal developed by the oscillator circuit 652A to the clock generator 654B of the sub-processor system 10B. For fault isolation, and to maintain signal quality, the M(underscore)CLK signal is delivered to the clock generator 654A of the sub-processor system 10A through a separate driver and a loopback connection 668. The reason for the... the cable (not shown) will establish the connection shown in Fig. 25 between the sub-processor systems 10A, 10B; connected another way, the connections will be similar, but the oscillator 652B... Fig. 25, the M(underscore)CLK signal produced by the oscillator circuit 652A of sub-processing system 10A is used by both sub-processing systems 10A, 10B as their respective SYNC CLK signals and the various other clock signals... produced by the clock generators 654A, 654B. Thereby, the clock signals of the paired sub-processing systems 10A, 10B are synchronized for the frequency locked operation necessary for duplex mode.

The VCXOs 662 of the clock... This allows both clock generators 654A, 654B to continue to provide to the two sub-processing systems 10A, 10B clock signals in the face of improper operation of the oscillator circuit 652A, although the sub-processor systems may no longer be frequency-locked.

The LOCK signals asserted by the phase comparators... LOCK signal signifies that the 50 Mhz signals produced by a clock generator 654 are synchronized, both in phase and in frequency, to the M(underscore)CLK signal. Thus, if either... signal that accompanies the symbol stream, and is used to push symbols onto the clock synchronizing FIFO of the receiving element (router 14, or CPU 12) is substantially identical in frequency... not phase, to that of the receiving element used to pull symbols from the clock synchronization FIFOs. For example, referring to Fig. 23, which illustrates symbols being sent from the router... clock (Local Clk). The former (Rcv Clk) is used to push symbols onto the clock synchronization FIFOs 126 of each CPU, whereas the latter is used to pull symbols from the... much higher frequency clock signal. In such situations provision must be made to ensure that synchronization is maintained between the two CPUs as to symbols pulled from the clock synchronization FIFOs 126 of each.

Here, a constant ratio clocking mechanism is used to control operation of the two clock synchronization FIFOs 126, providing the clock signal that pulls symbols from the two FIFOs at the... control mechanism is shown, designated with the reference numeral 70. As Fig. 26A illustrates, clock synchronization FIFO control mechanism 700 includes an pre-settable, multi-stage serial shift register 702, the... ratio of the clock signal at which symbols are communicated and pushed onto the clock synchronization FIFOs 126 to the frequency of the clock signal used locally. Here, a 15 stages... that will be used as the Local Clk signal to pull symbols from the clock synchronization FIFOs 126, and to operate (update) the pull pointer counter 130. The selected output is... of the CPU 12 to the clock signal used to push symbols onto the clock synchronization FIFO 126, Rcv Clk, the serial shift register is preset so that M stages of... duplexed CPUs 12 with a 50 Mhz clock. Thus, symbols are pushed onto the clock synchronization FIFOs 126 of the CPUs at a 50 Mhz rate. Assume further that the clock... of the MUX 704, which produces the clock signal that pulls symbols from the clock synchronization FIFOs 126, Rcv Clk, will contain, for each 100 ns... five symbols will be pushed onto, and five symbols will be pulled from, the clock synchronization FIFOs 126.

This example is symbolically shown in Fig. 26B, while the timing diagram shown... labelled "IN" in Fig. 27) of the Rcv Clk will push symbols onto the clock synchronization FIFOs 126. During that same 100 ns period, the serial shift register 702 circulates a... clocks which would require additional storage (i.e., an increase in the size of the synchronization FIFO) and impose more latency.

The constant ratio clock circuit presented here (Figs. 26) is... frequency to a clock regime of a different, higher frequency. The use of a clock synchronization FIFO is necessary here for compensating effects of signal delays when operating in synchronized, duplexed mode to receive pairs of identical command/data symbols from two different sources. However... so long as there are at least two registers in the place of the clock synchronization FIFO. Transferring data from a higher-frequency clock regime to a lower frequency clock regime... a wide range of possible clock ratios.

I/O Packet Interface:

Each of the sub-processor systems 10A, 10B, etc. will have some input/output capability, implemented with various peripheral units, although it is conceivable that the I/O of other sub-processor systems would be available so that a sub-processing system may not necessarily have local I/O. In any event, if local I/O... device (e.g., a signal line) would be received by the I/O packet interface unit 16 and used to form an interrupt packet that is sent to the CPU 12... OLAP bus, configuration information.

On-Line Access Port:

The MP 18 connects to the interface unit 24, memory controller (MC) 26, routers 14, and I/O packet interfaces with interface signals... OLAP 258 is essentially the same, regardless of what element (e.g. router 14, interface unit 24, etc.) it is used with. Fig. 28 diagrammatically illustrates the general structure of the... circuit chip used to implement certain of the elements discussed herein. For example, each interface

unit 24, memory controller 26, and router 14 is implemented by an application specific integrated circuit... of the OLAP 158 shown in Fig. 28 describes the OLAP associated with the interface unit 24, the MC 26, and the router 14 of the system.

As Fig. 28 shows... asymmetric variables, a "soft-vote" (SV) logic element 900 (Fig. 30A) is provided each interface unit 24 of each CPU 12. As Fig. 30 illustrates, the SV logic elements 900 of each interface unit 24 are connected to one another by a 2-bit SV bus 902, comprising bus lines 902a and 902b. Bus lines 902a carry one-bit values from the interface units 24 of CPU 12A to those of CPU 12B. Conversely, bus line 902b carries one... the CPU 12A.

Illustrated in Fig. 30B, is the SV logic element 900a of interface unit 24a of CPU 12A. Each SV logic element 900 is substantially identical in construction and... 900a should be understood as applying equally to the other logic elements 900a (of interface unit 24b, CPU 12A), and 900b (of the interface units 24a, 24b of CPU 12B) unless noted otherwise. As Fig. 30B illustrates, the SV logic... the logic elements 900a (as well as its own). In this manner the two interface units 24a, 24b of the CPU 12A can communicate asymmetrical variables to each other.

In a... to the remote register 907 of logic element 902a (and that of the other interface unit 24b).

The logic elements 902 form a part of the configuration registers 74 (Fig. 5). Thus, they may be written by the processor unit(s) 20 by communicating the necessary data/address information over at least a portion of... local and remote registers 906 and 907.

The MUX 914 operates to provide each interface unit 24 of CPU 12A with selective use of the bus line 902a for the SV logic elements 900a, or for communicating a BUS ERROR signal if encountered during the reintegration process (described below) used to bring a pair of CPUs 12 into lock-step, duplex operation... same time, write the enable registers 912 of the logic element 900 of both interface units 24 of each CPU. One of the two logic elements 900 of each CPU will... it is the output enable registers 912 associated with the logic elements 900 of interface units 24a of both CPUs 12A, 12B that are written to enable the associated drivers 916. Thus, the output registers 904 of the interface units 24a of each CPU will be communicated to the bus lines 902; that is, the... to the bus line 902a, while the output register associated with logic element 900b, interface unit 24a of CPU 12B is communicated to bus line 902b. The CPUs 12 will both... again written by each CPU, followed again by reading the remote input registers 907. This process is repeated, one bit at a time, until the entire variable is communicated from the... each CPU 12 to the remote input register of the other. Note that both interface units 24 of CPU 12B will receive the bit of asymmetric information.

One example of use... elements 900 are also used to communicate bus errors that may occur during the reintegration process to be described. When reintegration is being conducted, a REINT signal will be asserted. As... ERROR signal is selected by the MUX 914 and communicated to the bus line 902a.

Synchronization:

Proper operation of the sub-processing systems 10A, 10B (Figs. 1A, 2) whether operating independently (simplex mode), or paired and operating in synchronized lock-step (duplex mode), requires assurance that data communicated between the CPUs 12A, 12B and... routers 14A, 14B will be received properly, and that any initial content of the clock synchronization FIFOs 102 (of CPUs 12A, 12B; Fig. 5) and 519 (of routers 14A, 14B; Fig... erroneously interpreted as data or commands. The push and pull pointers of the various clock synchronization

FIFOs 102 (in the CPUs 12) and 518 (in the routers 14) need to be... apart, and presetting the associated FIFO queues to some known state. This done, all clock synchronization FIFOs are initialized for near frequency operation. Thus, when the system 10 is initially brought... in order to properly implement the lock-step operation of duplex mode operation, the clock synchronization FIFOs must be synchronized to operate with the particular source from which they receive data in order accommodate any... 14A, 14B to the CPUs 12A, 12B must be accounted for. It is the clock synchronization FIFOs 102 of the paired CPUs 12 that operate to receive message packet symbols, adjust... and present symbols to the two CPUs in a simultaneous manner to maintain lock-step synchronization necessary for duplex mode operation.

In similar fashion, each symbol received by the routers 14A... the CPUs (which is discussed further hereinafter). Again, it is the function of the clock synchronization FIFOs 518 of the routers 14A, 14B that receive message packets from the CPUs 12... so that the symbols received from the two CPUs 12 are retrieved from the clock synchronization FIFOs simultaneously.

Before discussing how the clock synchronization FIFOs of the CPUs and routers are reset, initialized, and synchronized, an understanding of their operation to maintain synchronous lock-step duplex mode operation is believed helpful. Thus, referring for the moment to Fig. 23, the clock synchronization FIFOs 102 of the CPUs 12A, 12B that receive data, for example, from the router... underscore)Clk, from the router 14A to the CPU 12B.

Consider operation of the clock synchronization FIFOs 102(sub(x)), 102(sub(y)), to receive identical symbol streams during duplex operation... held by the push and pull pointer counters 128, 130 for the CPU 12A (interface unit 24a), and the content of each of the four storage locations (byte 0. ..., byte 3... 6 show the same thing for the FIFO 102(sub(y)) of CPU 12B interface unit 24a for each symbol of the duplicated symbol stream.

Assuming the delay 640 is no... 0" locations of the queues 126. This is because (1) the FIFOs 102 have been synchronized to operate in synchronism (a process described below), and (2) the push pointer counters 128 are clocked by the clock signal... of the symbol stream transmitted by the router 14A will be pulled from the clock synchronization FIFOs 102 of the CPUs 12A, 12B simultaneously, maintaining the required synchronization of received data when operating in duplex mode. In effect, the depths of the queues... order to achieve the operation just described with reference to Table 6, the reset and synchronization process shown in Fig 31A is used. The process not only initializes the clock synchronization FIFOs 102 of the CPUs 12A, 12B for duplex mode operation, but also operates to adjust the clock synchronization FIFOs 518 (Fig. 19A) of the CPU ports of each of the routers 14A, 14B for duplex operation. The reset and synchronization process uses the SYNC command symbol to initiate a time period, delineated by the SYNC CLK signal 970 (Fig. 31B), to reset and initialize the respective clock synchronization FIFOs of the CPUs 12A and 12B and routers 14A, 14B. (me SYNC CLK signal... It is of a lower frequency than that used to receive symbols by the clock synchronization FIFOs, T(underscore)Clk. For example, where T(underscore)Clk is approximately 50 MHz, the... signal is approximately 3.125 MHz.)

Turning now to Fig. 31A, the reset and initialization process begins at step 950 by switching the clock signals used by the CPUs 12A, 12B and routers 14A, 14B as the transmit (T(underscore)Clk) and the unit's local clock (Local Clk) clock signals so that they are derived from the same... In addition, configuration registers in the CPUs 12A, 12B (configuration registers 74 in the interface units 24) and the routers 14A, 14B (contained in control logic unit 509 of routers 14A, 14B) are set to the FreqLock state.

The following discussion involves step 952, and makes reference to the interface unit 24 (Fig.5), router 14A (Fig. 19A) and Figs. 31A and 31B. With the clock... otherwise be sent followed by a self-addressed message packet. Any message packet in the process of being received and retransmitted when the SLEEP command symbols are received and recognized by... per the destination address). The SLEEP command symbol operates to "quieece" router 14A for the synchronization process. The self-addressed message packet sent by the CPU 12A, when received back by the... message packet sent after the SLEEP command symbol would necessarily have to be the last processed by the router 14A.

At step 954 the CPU 12A checks to see if it... ..the router will assert a RESET signal 972 that is applied to the two clock synchronization FIFOs 518 contained in the input logic 505(sub(4)), 505(sub(5)) of the... ..receive symbols directly from CPUs 12A, 12B. RESET, while asserted, will hold the two clock synchronization FIFOs 518 in a temporarily non-operating reset state with the push and pull pointer... ..As each of the CPUs 12 receive SYNC symbols are detected by the storage and processing units of the packet receivers 96 (Figs. 5 and 6) cause the RESET signal to be asserted by the packet receivers 96 (actually, storage and processing elements 110; Fig. 6) of each CPU 12. the RESET signal is applied to the... ..4))), CPUs 12 and routers 14A, 14B de-assert the RESET signals, and the clock synchronization FIFOs of the CPUs 12A, 12, and...the delay, the router 14A and CPUs 12 resume pulling data from their respective clock synchronization FIFOs and resume normal operation. The clock synchronization FIFOs of the router 14A begin pulling symbols from the queue (previously set by RESET... ..from the CPU 12A with the T(underscore)Clk will be pushed onto the clock synchronization FIFO at, for example, queue location 0 (or whatever other location pointed to by the... ..0 (or whatever other location the push pointer was set to by RESET). The clock synchronization FIFOs of the router 14A are now synchronized to accommodate whatever delay 640 may be present in one communications path, relative to the... ..and the CPUs 12A, 12B.

Similarly, at the same virtual time, operation of the clock synchronization FIFOs 102 of both CPUs 12A, 12B is resumed, synchronizing them to the router 14A. Also, the CPUs 12A, 12B quit sending the SLEEP command... ..in favor of READY symbols, and resume message packet transmission, as appropriate.

That completes the synchronization process for the router 14A. However, the process must also be performed for the router 14B. Thus, the CPU 12A returns to step... ..however, assuming that the CPUs 12A, 12B are operating in duplex mode, the method and apparatus used to detect and handle a possible error, resulting in divergence of the CPUs from... ..via a message packet destined for a peripheral device of one or the other sub-processor systems 10A, 10B. Depending upon the destination of the outgoing message packet, step 1002 will... ..router 14 will issue an ERROR signal to the router control logic 509, causing the process to move to step 1004 where the router 14 detecting divergence will transmit a DVRG... ..time outs to occur. A router detecting divergence (without also detecting any simple link error) buys itself time to check the CRC of the received message packet by waiting for the... ..router 14, or received, all further message packets received from the CPUs and in the process of being routed when divergence was detected, or the DVRG symbol received, will be passed... ..1010) contained in a one of the configuration registers 74 (Fig. 5) of the interface unit 24 of each CPU.

Returning for the moment to step 1006, the determination of which... ..local" is meant to refer to the router 14A, 14B contained in the same sub-processor system 10A, 10B as the CPU. For example, referring to Fig. 1A, router 14A is... ..bit mentioned above: the bit contained in one of the configuration registers 74 of interface unit 24 (Fig. 5) of each CPU. When set to a first state, that particular CPU... ..the other CPU. In response, the state machines (not shown) within the control and status unit 509 (Fig. 19A) changes the "favorite" bits described above.

A few examples may facilitate understanding... ..DVRG symbol will echo that symbol to the routers 14A, 14B, start its internal divergence process timer, and begin determination of whether to continue or terminate. Having received a TLB symbol... ..to diverge with no errors reported. This can happen only if software (running on the processors 20) uses known divergent data to alter state. For example, suppose each CPU 12 has... ..number of the CPU 12A will differ from that of the CPU 12B. If the processors use the serial number to change the sequence of instructions executed (say, by branching if the serial number comes after some value) or to modify the value contained in a processor register, the complete "state" of the CPUs 12 will differ. In such cases, the "asymmetrical... ..of the primary CPU simply allows one CPU, and thereby the system 10, to continue processing without software intervention.

- An error at the output of the interface unit 24 of a CPU 12 will be detected by the router 14A, 14B, depending upon... ..router 14A, 14B that connects to a CPU 12 will be detected by the interface unit 24 of the affected CPU.

The CPU will send a TLB symbol to the faulty... possible failure and, without external intervention, and transparently to the system user, remove the failing unit (CPU 12A or 12B, or router 14A or 14B) from the system to obviate or... reintegration." The discussion will refer to the CPUs 12A, 12B, routers 14A, 14B, and maintenance processor 18A, 18B shown forming parts of the processing system 10 illustrated in Fig. 1A. In addition, discussion will refer to the processors 20a, 20b, the interface units 24a, 24b, and the memory controllers 26a, 26b (Fig. 2) of the CPUs 12A, 12B as single units, since that is the way they function.

Reintegration is used to place two CPUs in... both of the paired CPUs at virtually the same time.

The major steps in the process for changing from simplex mode operation of the one on-line CPU to duplex mode... greater detail by the flow diagrams of Figs. 33A - 33D, generally are:

1. Setup and synchronize the two CPUs (one on-line, the other off-line) and their connected routers to... the memory of the on-line CPU to the off-line CPU, maintaining a tracking process that monitors changes in the memory of the on-line CPU that have not been... and may need to be copied over to, the off-line CPU;
3. Setup and synchronize the CPUs to run a delayed (slave) duplex mode from the same instruction stream (lock... will write the predetermined registers (not shown) of the control registers 74 in the interface units 24 of CPUs 12A and 12B, to a next state (after a soft operation) in... the off-line CPU 12B.

Next, a sequence is entered (steps 1060 - 1070) that will synchronize the clock synchronization FIFOs of the CPUs 12A, 12B and routers 14A, 14B in much the same fashion... the same steps described above in connection with the discussion of Figs. 31A, 31B to synchronize the clock synchronization FIFOs. The on-line CPU 12A will send the sequence of a SLEEP symbol, self-addressed message packet, and SYNC symbol which, with the SYNC CLK signal, operates to synchronize CPUs and routers. Once so synchronized, the on-line CPU 12A then, at step 1066, sends a Soft Reset (SRST) command... of all configuration registers and control registers (e.g., configuration registers 74 of the interface units 24) cache, and the like to memory 28 of the on-line CPU, copying the... time to have the system 10 off-line for reintegration. For that reason, the reintegration process is performed in a manner that allows the on...not match that of the off-line CPU. The reason for this is that normal processing by the processor 20 of the on-line CPU can change memory content after it has been copied... when a memory location is written in the on-line CPU 12A during the reintegration process it is marked as "dirty;" second, all copying of memory to the off-line CPU... may, however, limit the ability to detect two-bit errors. But, since the memory copying process will last for a only relatively short period of time, this risk is believed acceptable... memory location in CPU 12A is made (either an incoming I/O write, or a processor write operation). The returning data (that was copied over to the off-line CPU) would... controller 26 (Fig. 2) of the on-line CPU to monitor memory locations in the process of being copied over to the off-line CPU 12B. The memory controller uses a... within the block had been written by another operation (e.g., a write by the processor 20, an I/O write, etc.), that prior write operation will flag the location in... still must be copied over to the off-line CPU 12B.

Returning to the reintegration process, and now to Fig. 33B, the memory tracking (AtomicWrite mechanism and using ECC to mark... entails writing a reintegration register (not shown; one of the configuration registers 74 of interface unit 24 - Fig. 5) to cause a reintegration (REINT) signal to be asserted. The REINT signal... is left alone. Throughout the incremental copy operations, the normal actions of the on-line processor will mark some memory locations dirty.

Several passes of incremental copying will need to be... the number of successful WriteConditional operations at the end of each pass through memory, the processors 20 can determine the effect of a given pass compared to the previous pass. When the benefits drop off, the processors 20 will give up on the precopy operations. At this point the reintegration process is ready to place the two CPUs 12A, 12B into lock-step operation.

Thus, the... in Fig. 33C, where at step 1100, the on-line CPU 12A momentarily halts foreground processing, i.e., execution of a user application. The remaining state (e.g., configuration registers, cache, etc.) of the on-line

processors 20 and its caches is then read and written to a buffer (series of memory... ..to the off-line CPU 12B, together with a "reset vector" that will direct the processor units 20 of both CPUs 12A, 12B to a reset instruction.

Next, step 1106 will quiesce... ..to ensure that the FIFOs of the routers are clear, that the FIFOs of the processor interfaces 24 are clear, and no further incoming I/O message packets are forthcoming. At... ..symbol will be received and acted upon by both CPUs 12A, 12B, to cause the processor units 20 of each CPU to jump to the location in memory 28 containing the reset... ..a subroutine that will restore the stored state of both CPUs 12A, 12B to the processor units 20, caches 22, registers, etc. The CPUs 12A, 12B will then begin executing the same... ..enabling of the ECC bit to mark dirty locations must now be disabled, since the processors are doing the same thing to the same memory. During this stage of the reintegration... ..encountered by CPU 12A.

Meanwhile, the bus error in the CPU 12A will cause the processor unit 20 to be forced into an error-handling routine to determine (1) the cause of... ..error was caused by an attempt to read a memory location marked dirty. Accordingly, the processor unit 20 will initiate (via the BTE 88 -- Fig. 5) the AtomicWrite mechanism to copy the... ..the SRST symbols are now received by the CPUs 12A, 12B, they will cause both processor units 20 of the CPUs to be reset to start from the same location with the... ..will periodically update, e.g., a database or audit file that is indicative of the processing of the primary CPU up to that point in time of the update. Should the... ..in error-checking redundancy to the CPU 12B, in the same manner that the individual processor units 20a, 20b of the CPU 12A provide fail-fast, fault tolerance for the CPU - when... ..cost system is applicable, as illustrated in Fig. 34. As shown in Fig. 34, a processing system 10' includes the CPU 12A and routers 14A, 14B structured as described above. The... ..and the CPUs are also the same.

Thus, the CPU 12B' comprises only a single processor unit 20' and associated support components, including the cache 22', interface unit (IU) 24', memory controller 26', and memory 28'. Thus, while the CPU 12A is structured in the manner shown in Fig. 2, with cache processor unit, interface unit, and memory control redundancies, approximately one-half of those components are needed to implement CPU... ..stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... ..i.e, parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of... ..inputs connected to receive the data output

from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... ..for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the... ..attached, without the assistance of any other CPU 12 in the system. Many prior parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O device becomes... ..rest of the system. Other prior systems provide access to I/O through pairs of processors so that should one of the processors fail, access to the corresponding I/O is still available through the remaining I/O... ..if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... ..multiprocessing system access to every peripheral, as done here, operates to extend the "primary"/"backup" process taught in the above-identified U.S. Patent No. 4,228,496. There, a multiple CPU system may have a primary process may running on one CPU, while a backup process resides in the background on another of the CPUs. Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which

the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be... ..e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system... ..CPU can access any I/O without the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... ..four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... ..an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time... ..indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That...When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... ..response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Specification: ...CPU 12B.

Next, a sequence is entered (steps 1060 - 1070) that will synchronize the clock synchronization FIFOs of the CPUs 12A, 12B and routers 14A, 14B in much the same fashion... ..the same steps described above in connection with the discussion of Figs. 31A, 31B to synchronize the clock synchronization FIFOs. The on-line CPU 12A will send the sequence of a SLEEP symbol, self-addressed message packet, and SYNC symbol which, with the SYNC CLK signal, operates to synchronize CPUs and routers. Once so synchronized, the on-line CPU 12A then, at step 1066, sends a Soft Reset (SRST) command... ..of all configuration registers and control registers (e.g., configuration registers 74 of the interface units 24) cache, and the like to memory 28 of the on-line CPU, copying the... ..time to have the system 10 off-line for reintegration. For that reason, the reintegration process is performed in a manner that allows the on-line CPU to continue executing user... ..not match that of the off-line CPU. The reason for this is that normal processing by the processor 20 of the on-line CPU can change memory content after it has been copied... ..when a memory location is written in the on-line CPU 12A during the reintegration process it is marked as "dirty;" second, all copying of memory to the off-line CPU... ..may, however, limit the ability to detect two-bit errors. But, since the memory copying process will last for a only relatively short period of time, this risk is believed acceptable...memory location in CPU 12A is made (either an incoming I/O write, or a processor write operation). The returning data (that was copied over to the off-line CPU) would... ..controller 26 (Fig. 2) of the on-line CPU to monitor memory locations in the process of being copied over to the off-line CPU 12B. The memory controller uses a... ..within the block had been written by another operation (e.g., a write by the processor 20, an I/O write, etc.), that prior write operation will flag the location in... ..still must be copied over to the off-line CPU 12B.

Returning to the reintegration process, and now to Fig. 33B, the memory tracking (AtomicWrite mechanism and using ECC to mark... ..entails writing a reintegration register (not shown; one of the configuration registers 74 of interface unit 24 - Fig. 5) to cause a reintegration (REINT) signal to be asserted. The REINT signal... ..is left alone. Throughout the incremental copy operations, the normal actions of the on-line processor will mark some memory locations dirty.

Several passes of incremental copying will need to be... ..the number of successful WriteConditional operations at the end of each pass through memory, the processors 20 can determine the effect of a given pass compared to the

previous pass. When the benefits drop off, the processors 20 will give up on the precopy operations. At this point the reintegration process is ready to place the two CPUs 12A, 12B into lock-step operation.

Thus, the... ..in Fig. 33C, where at step 1100, the on-line CPU 12A momentarily halts foreground processing, i.e., execution of a user application. The remaining state (e.g., configuration registers, cache, etc.) of the on-line processors 20 and its caches is then read and written to a buffer (series of memory... ..to the off-line CPU 12B, together with a "reset vector" that will direct the processor units 20 of both CPUs 12A, 12B to a reset instruction.

Next, step 1106 will quiesce... ..to ensure that the FIFOs of the routers are clear, that the FIFOs of the processor interfaces 24 are clear, and no further incoming I/O message packets are forthcoming. At... ..symbol will be received and acted upon by both CPUs 12A, 12B, to cause the processor units 20 of each CPU to jump to the location in memory 28 containing the reset... ..a subroutine that will restore the stored state of both CPUs 12A, 12B to the processor units 20, caches 22, registers, etc. The CPUs 12A, 12B will then begin executing the same... ..enabling of the ECC bit to mark dirty locations must now be disabled, since the processors are doing the same thing to the same memory. During this stage of the reintegration... ..encountered by CPU 12A.

Meanwhile, the bus error in the CPU 12A will cause the processor unit 20 to be forced into an error-handling routine to determine (1) the cause of... ..error was caused by an attempt to read a memory location marked dirty. Accordingly, the processor unit 20 will initiate (via the BTE 88 -- Fig. 5) the AtomicWrite mechanism to copy the... ..the SRST symbols are now received by the CPUs 12A, 12B, they will cause both processor units 20 of the CPUs to be reset to start from the same location with the... ..will periodically update, e.g., a database or audit file that is indicative of the processing of the primary CPU up to that point in time of the update. Should the... ..in error-checking redundancy to the CPU 12B, in the same manner that the individual processor units 20a, 20b of the CPU 12A provide fail-fast, fault tolerance for the CPU - when... ..cost system is applicable, as illustrated in Fig. 34. As shown in Fig. 34, a processing system 10' includes the CPU 12A and routers 14A, 14B structured as described above. The... ..and the CPUs are also the same.

Thus, the CPU 12B' comprises only a single processor unit 20' and associated support components, including the cache 22', interface unit (IU) 24', memory controller 26', and memory 28'. Thus, while the CPU 12A is structured in the manner shown in Fig. 2, with cache

processor unit, interface unit, and memory control redundancies, approximately one-half of those components are needed to implement CPU... ..stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... ..i.e., parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of... ..inputs connected to receive the data output

from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... ..for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the... ..attached, without the assistance of any other CPU 12 in the system. Many prior parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O device becomes... ..rest of the system. Other prior systems provide access to I/O through pairs of processors so that should one of the processors fail, access to the corresponding I/O is still available through the remaining I/O... ..if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... ..multiprocessing system access to every peripheral , as done here, operates to extend the "primary"/"backup" process taught in the above-identified U.S. Patent No. 4,228,496. There, a multiple CPU system may have a primary process may running on one CPU, while a backup process resides in the background on another of the CPUs. Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be... ..e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system... ..CPU can access any I/O without the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... ..four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... ..an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time... ..indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... ..When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... ..response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Claims: ...A2

1. Apparatus for providing synchronized clock signals to at least a pair of synchronous processing elements, comprising:

each of the pair of processing elements including:

a master oscillator circuit to produce a master clock signal,

a clock generator circuit coupled to receive the master clock signal to produce the synchronized clock signals therefrom, the clock generator circuit having a voltage controlled oscillator circuit responsive to... ..first clock signal produce a number of divisions of the first clock signal forming the synchronized clock signals and a replica of the master clock signal, and a phase compare means... ..and means for coupling the master oscillator circuit of a one of the pair of processing elements to provide the master clock signal for the clock generator to the pair of processing elements.

6/K/26 (Item 26 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Patent Assignee:

- Compaq Computer Corporation... ;
;;

Country	Number	Kind	Date	
Type		Pub. Date	Kind	Text
...Compaq Computer Corporation (687797) 20555 SH 249 Houston, Texas 77070-2698 US		19		
Examination...		19		
Available Text		Language	Update	Word Count
Total Word Count (Document A)				
Total Word Count (Document B)				
Total Word Count (All Documents)				

Specification: ...and 08/485,055 filed concurrently herewith.

The present invention is directed generally to data processing systems, and more particularly to a multiple processing system and a reliable system area network that provides connectivity for interprocessor and input/output... ..and communications systems to general purpose high availability commercial systems. The evolution of fault tolerant computers has been well documented (see D. P. Siewiorek, R. S. Swarz, "The Theory and Practice... ..and the Jet Propulsion laboratory began to apply fault tolerance to the development of guidance computers for aerospace applications. The 1960's also saw the development of the first AT&T electronic switching systems.

The first commercial fault tolerant machines were introduced by Tandem Computers in the 1970's for use in on-line transaction processing applications (J. Bartlett, "A NonStop Kernal," in proc. Eighth Symposium on Operating System Principles, pp... ..systems were introduced in the 1980's (O. Serlin, "Fault- Tolerant Systems in Commercial Applications," Computer, pp. 19-30, August 1984). Current commercial fault tolerant systems include distributed memory multi-processors, shared-memory transaction based systems, "pair-and-spare" hardware fault tolerant systems (see R. Freiburghouse, "Making Processing Fail-safe," Mini-micro Systems, pp. 255-264, May 1982; U.S. Patent No. 4... ..system.), and triple-modular-redundant systems such as the "Integrity" computing system manufactured by Tandem Computers Incorporated of Cupertino, California, assignee of this application and the invention disclosed herein.

Most applications of commercial fault tolerant computers fall into the category of on-line transaction processing. Financial institutions require high availability for electronic funds transfer, control of automatic teller machines, and... ..telecommunications systems.

Vendors of fault tolerant machines attempt to achieve both increased system availability, continuous processing, and correctness of data even in the presence of faults. Depending upon the particular system architecture, application software ("processes") running on the system either continue to run despite failures, or the processes are automatically restarted from a recent checkpoint when a fault is encountered. Some fault tolerant systems are provided with sufficient component redundancy to be able reconfigure around failed components, but processes running in the failed modules are lost. Vendors of commercial fault tolerant systems have extended fault tolerance beyond the processors and disks. To make large improvements in reliability, all sources of failure must be addressed... ..power supplies, fans and inter-module connections.

The "NonStop," and "Integrity" architectures manufactured by Tandem Computers Incorporated, (both respectively illustrated broadly in U.S. Patent No. 4,228,496 and U... ..assigned to the assignee of this application; NonStop and

Integrity are registered trademarks of Tandem Computers Incorporated) represent two current approaches to commercial fault tolerant computing. The NonStop system, as generally... above-identified U.S. Patent No. 4,278,496, employs an architecture that uses multiple processor systems designed to continue operation despite the failure of any single hardware component. In normal operation, each processor system uses its major components independently and concurrently, rather than as "hot backups". The NonStop system architecture may consist of up to 16 processor systems interconnected by a bus for interprocessor communication. Each processor system has its own memory which contains a copy of a message-based operating system. Each processor system controls one or more input/output (I/O) busses. Dual-porting of I/O controllers and devices provides multiple paths to each device. External storage (to the processor system), such as disk storage, may be mirrored to maintain redundant permanent data storage.

This... hardware, while fault recovery is the responsibility of the software.

Also, in the Nonstop multi-processor architecture, application software ("process") may run on the system under the operating system as "process-pairs," including a primary process and a backup process. The primary process runs on one of the multiple processors while the backup process runs on a different processor. The backup process is usually dormant, but periodically updates its state in response to checkpoint messages from the primary process. The content of a checkpoint message can take the form of complete state update, or currently most application code runs under transaction processing software which provides recovery through a combination of checkpoints and transaction two-phase commit protocols.

Interprocessor message traffic in the Tandem Nonstop architecture includes each processor periodically broadcasting an "I'm Alive" message for receipt by all the processors of the system, including itself, informing the other processors that the broadcasting processor is still functioning. When a processor fails, that failure will be announced and identified by the absence of the failed processor's periodic "I'm Alive" message. In response, the operating system will direct the appropriate backup processes to begin primary execution from the last checkpoint. New backup processes may be started in another processor, or the process may be run with no backup until the hardware has been repaired. U.S. Patent... example of this technique.

Each I/O controller is managed by one of the two processors to which it is attached. Management of the controller is periodically switched between the processors. If the managing processor fails, ownership of the controller is automatically switched to the other processor. If the controller fails, access to the data is maintained through another controller.

In addition to providing hardware fault tolerance, the processor pairs of the above-described architecture provide some measure of software fault tolerance. When a processor fails due to a software error, the backup processor frequently is able to successfully continue processing without encountering the same error. The software environment in the backup processor typically has different queue lengths, table sizes, and process mixes. Since most of the software bugs escaping the software quality assurance tests involve infrequent data dependent boundary conditions, the backup processes often succeed.

In contrast to the above-described architecture, the Integrity system illustrates another approach... fault recovery is the logical choice since few modifications to the software are required. The processors and local memories are configured using triple-modular-redundancy (TMR). All processors run the same code stream, but clocking of each module is independent to provide tolerance... three streams is asynchronous, and may drift several clock periods apart. The streams are re-synchronized periodically and during access of global memory. Voters on the TMR Controller boards detect and mask failures in a processor module. Memory is partitioned between the local memory on the triplicated processor boards and the global memory on the duplicated TMRC boards. The duplicated portions of the... techniques to detect failures. Each global memory is dual ported and is interfaced to the processors as well to the I/O Processors (IOPs). Standard VME peripheral controllers are interfaced to a pair of busses through a Bus...

...the BIMs to switch control of all controllers to the remaining IOP. Mirrored disk storage units may be attached to two different VME controllers.

In the Integrity system all hardware failures... ..reintegrated on-line.

The preceding examples illustrate present approaches to incorporating fault tolerance into data processing systems. Approaches involving software recovery require less redundant hardware, and offer the potential for some... ..have been developed on other systems.

Thus, the systems described above provide fault tolerant data processing either by hardware (e.g, fail-functional, employing redundancy) or by software techniques (fail-fast... ..hardware). However, none of the systems described are believed capable of providing fault tolerant data processing, using both hardware (fail-functional) and software (fail-fast) approaches, by a single data processing system.

Computing systems, such as those described above, are often used for electronic commerce: electronic data interchange (EDI) and global messaging. Today's demands upon such electronic commerce, however, is demanding more and more throughput capacity as the number of users increases and... ..networks such as local area networks (LANS), and the like.

A key requirement for a server architecture is the ability to move massive quantities of data. The server should have high bandwidth that is scalable, so that added throughput capacity can be added... ..response time, latency affects service levels and employee productivity.

The present invention provides a multiple-processor system that combines both of the two above-described approaches to fault tolerant architecture, hardware redundancy and software recovery techniques, in a single system.

Broadly, the present invention includes a processing system composed of multiple sub-processing systems. Each sub-processing system has, as the main processing element, a central processing unit (CPU) that in turn comprises a pair of processors operating in lock-step, synchronized fashion ...execute each instruction of an instruction stream at the same time. Each of the sub-processing systems further include an input/output (I/O) system area network system that provides redundant communication paths between various components of the larger processing

system, including a CPU and assorted peripheral devices (e.g., mass storage

units, printers, and the like) of a sub-processing system, as well as between the sub-processors that may make up the larger overall processing system. Communication between any component of the processing system (e.g., a CPU and another CPU, or a CPU and any peripheral device, regardless of which sub-processing system it may belong to) is implemented by forming and transmitting packetized messages that are... ..responsible for choosing the proper or available communication paths from a transmitting component of the processing system to a destination component based upon information contained in the message packet. Thus, the... ..peripherals, but permits it to also be used for interprocessor communications.

As indicated above, the processing system of the present invention is structured to provide fault-tolerant operation through both "fail... ..at a variety of points in the various data paths between the (lock-step operated) processor elements of the CPU and its associated memory. In particular, the processing system of the present invention conducts error-checking at an interface, and in a manner... ..little impact on performance. Prior art systems typically implement error-checking by running pairs of processors, and checking (comparing) the data and instruction flow between the processors and a cache memory. This technique of error-checking tended to add delay to the... ..error-checking precluded use of off-the-shelf parts that may be available (i.e., processor/cache memory combinations on a single semiconductor chip or module). The present invention performs error-checking of the processors at points that operate at slower rates, such as the main memory and I/O interfaces which operate at slower speeds than the processor-cache interface. In addition, the error-checking is performed at locations that allow detection of errors that

may occur in the processors, their cache memory, and the I/O and memory interfaces. This allows simpler designs for... other data integrity checks.

Error-checking of the communication flow between the components of the processing system is achieved by adding a cyclic-redundancy-check (CRC) to the message packets that... "Good" (TPG) or "This Packet Bad" (TPB) - is appended to every packet. A maintenance diagnostic processor can use this information to isolate a link or router element that introduces an error... of topologies, so that alternate paths can be provided between any two elements of a processing system (e.g., between a CPU and an I/O device), for communication in the... so (e.g., by creating a "deadlock" condition, discussed further below).

The CPUs of a processing system are capable of operating in one of two basic modes: a "simplex mode" in... independently of the other, or a "duplex" mode in which pairs of CPUs operate in synchronized, lock-step fashion. Simplex mode operation provides the capability of recovering from faults that are... U.S. Pat. No. 4,228,496 which teaches a multiprocessing system in which each processor has the capability of checking on the operability of its sibling processors, and of taking over the processing of a processor found or believed to have failed). When operating in duplex mode, the paired CPUs both... fault tolerant platform for less robust operating systems (e.g., the UNIX operating system). The processing system of the present invention, with the paired, lock-step CPUs, is structured so that... masked (i.e., operating despite the existence of a fault), primarily through hardware.

When the processing system is operating in duplex mode, each CPU pair uses the I/O system to access any peripheral of the processing system, regardless of which (of the two, or more) sub-processor system the peripheral may be ostensibly a member of. Also, in duplex mode, message packets... message for the CPU pair (from either a peripheral device such as a mass storage unit or from a processing unit), will replicate the message and deliver it to both CPUs of the pair using synchronization methods that ensure that the CPUs remain synchronized. In effect, the duplex CPU pair, as viewed from the I/O system and other... as a single CPU. Thus, the I/O system, which includes elements from all sub-processing systems, is made to be seen by the duplex CPU pair as one homogeneous system... a multiprocessor system in which the CPU of any one is actually a pair of synchronized, lock-step CPUs.

Yet another important aspect of the present invention is that interrupts issuing... interrupts via the message packet system ensures that they will arrive at duplexed CPUs in synchronized fashion, in the same manner as I/O message packets. Interrupt message packets will contain... the system. In addition, using the same messaging system to communicate data between I/O units and the CPUs and to communicate interrupts to the CPUs preserves the ordering of I... the implementation of a technique of validating access to the memory of any CPU. The processing system, as structured according to the present invention, permits the memory of any CPU to a CPU and any other component of the processor system. Thereby, the individual processor units of the CPU are removed from the more mundane tasks of getting information from memory and out onto the TNet network, or accepting information from the network. The processor unit of the CPU merely sets up data structures in memory containing the data to be... is required, where in memory the response is to be placed when received. When the processor unit completes the task of creating the data structure, the block transfer engine is notified to... response is received, it is routed to the expected memory location identified, and notifies the processor unit that the response was received.

Further aspects and features of the present invention will become... invention, which should be taken in conjunction with the accompanying drawings.

Fig. 1A illustrates a processing system constructed in accordance with the teachings of the present invention, and Figs. 1B and 1C illustrate two alternate configurations of the processing system of Fig. 1A, employing clusters or arrangements of the processing system of Fig. 1A;

Fig. 2 illustrates, in simplified block diagram form, the central processing unit (CPU) that forms a part of each sub-processor system of Figs. 1A - 1C;

Figs. 3A - 3D and 4A - 4C illustrate the construction of... ..the area network I/O system shown in Fig. 2;

Fig. 5 illustrates the interface unit that forms a part of the CPUs of Fig. 2 to interface the processor and memory with the I/O area network system;

Fig. 6 is a block diagram, illustrating a portion of packet receiver of the interface unit of Fig. 5;

Fig. 7A diagrammatically illustrates the clock synchronization FIFO (CS FIFO) used by the packet receiver section packet receiver shown in Fig. 6;

Fig. 7B is an block diagram of a construction of the clock synchronization FIFO structure shown in Fig. 7A;

Fig. 8 illustrates the cross-connections for error-checking outbound transmissions from the two interface units of a CPU;

Fig. 9 illustrates an encoded (8B to 9B) data/command symbol;

Fig. 10 illustrates the method and structure used by the interface unit of Fig. 5 to cross-check for errors data being transferred to the memory controllers... ..of a CPU of Fig. 2 to other (external to the CPU) components of the processing system;

Fig. 12 is a block diagram that diagrammatically illustrates the formation of an address... ..14A illustrates the logic for posting interrupt requests to queues in memory and to the processor units of the CPU of Fig. 2;

Fig. 14B illustrates the process used to form a memory address for a queue entry;

Fig. 15 is a block... ..data output constructs formed in the memory of the CPU of Fig. 2 by a processor unit, and containing data to be sent via the area I/O networks shown in Figs. 1A - 1C, and also illustrating the block transfer engine (BTE) unit of the interface unit of Fig. 5 that operates to access the data output constructs for transmission to the... ..pair of memory controllers between memory of a CPU of Fig. 2 and its interface unit for accessing from memory 72 bits of data, including two simultaneously-accessed 32-bit words... ..other for error-checking;

Fig. 19A is a simplified block diagram illustration of the router unit used in the area input/output networks of the processing systems shown in Figs. 1A - 1C;

Fig. 19B illustrates comparison on two port inputs of the router unit of Fig. 19A;

Fig. 20A is a block diagram the construction of one of the six input ports of the router unit shown in Fig. 19A;

Fig. 20B is a block diagram of the synchronization logic used to validate command/data symbols received at an input port of the router unit of Fig. 19A;

Fig. 21A is a block diagram illustration of the target port selection... ..is a block diagram illustration of one of the six output ports of the router unit shown in Fig. 19A;

Fig. 23 is an illustration of the method used to transmit identical information to a duplexed pair CPUs of Fig. 2 in synchronized fashion when the processing system is operating in lock-step (duplex) mode, using a pair the FIFOs of Fig... ..is a simplified block diagram illustrating the clock generation system of each of the sub-processing systems of Figs. 1A - 1C for developing the plurality of clock signals used to operate the various elements of that sub-processing system;

Fig. 25 illustrates the topology used to interconnect the clock generation systems of paired sub-processing systems for synchronizing the various clock signals of the pair of sub-processing systems to one another;

Fig. 26A and 26B illustrates a FIFO constant rate clock control logic used to control the clock synchronization FIFO of Figs. 8 or 20 in the situation when the two clocks used to... ..structure of the on-line access port (OLAP) used to provide access to the maintenance

processor (MP) to the various elements of the system of Fig. 1A (or those of Figs... ..the soft-flag logic used to handle asymmetric variables between the CPUs of paired sub-processor systems operating in duplex mode;

Fig. 31A shows a flow diagram, and Fig. 31B illustrates a portion of SYNC CLK, both of which are used to reset and synchronize the clock synchronization FIFOs of the CPUs and routers of the processing system of Fig. 1A that receive information from each other;

Fig. 32 is a flow... ..33A - 33D generally illustrate the procedure used to bring an one of the CPUs of processing system shown in Fig. 1A into lock-step, duplex mode operation with the other of the CPUs without measurably halting operation of the processing system; and

Fig. 34 illustrates a reduced cost architecture incorporating teachings of the invention; and... ..to the figures and, for the moment, principally Fig. 1A, there is illustrated a data processing system, designated with the reference 10, constructed according to the various teachings of the present invention. As Fig. 1A shows, the data processing system 10 comprises two sub-processor systems 10A and 10B each of which are substantially the same in structure and function... ..should be appreciated that, unless noted otherwise, a description of any one of the sub-processor systems 10 will apply equally to any other sub-processor system 10.

Continuing with Fig. 1A therefore, each of the sub-processor systems 10A, 10B is illustrated as including a central processing unit (CPU) 12, a router 14, and a plurality of input/output (I/O) packet interfaces... ..one of the I/O packet interfaces 16 will also have coupled thereto a maintenance processor (MP) 18.

The MP 18 of each sub-processor system 10A, 10B connects to each of the elements of that sub-processor system via an IEEE 1149.1 test bus 17 (shown in phantom in Fig. 1A... ..accompanying clock signal. As Fig. 1A further illustrates, TNet Links L also interconnect the sub-processor systems 10A and 10B to one another, providing each sub-processor system 10 with access to the I/O devices of the other as well as inter-CPU communication. As will be seen, any CPU 12 of the processing system 10 can be given access to the memory of any other CPU 12, although... ..the memory of a CPU 12 by a wayward peripheral device 17.

Preferably, the sub-processor systems 10A/10B are paired as illustrated in Fig. 1A (and Figs 1B and 1C, discussed below), and each sub-processor system 10A/10B pair (i.e., comprising a CPU 12, at least one router 14... ..12A) connects, by a TNet Link L to a router (14A) of the corresponding sub-processor system (e.g., 10A). Conversely, the Y port connects the CPU (12A) to the router (14B) of the companion sub-processor system (...for access by a CPU (12A) to the I/O devices of the other sub-processor system (10B), but also to the CPU (12B) of that system for inter-CPU communication.

Information is communicated between any element of the processing system 10 and any other element (e.g., CPU 12A of sub-processor system 10A) of the system and any other element of the system (e.g., an I/O device associated with an I/O packet interface 16B of sub-processor system 10B) via message "packets." Each message packet is made up of a number of... ..this reason, a unique method of receiving the symbols at the receiver, using a clock synchronization first-in-first-out (CS FIFO) storage structure (described more fully below), has been developed... ..operation means just that: the frequencies of the clock signals of the transmitter and receiver units are locked, although not necessarily in phase. Frequency locked clock signals are used to transmit symbols between the routers 14A, 14B and the CPUs 12 of paired sub-processor systems (e.g., sub-processor systems 10A, 10B, Fig. 1A). Since the clocks of the transmitting and receiving element are not phase related, a clock synchronization FIFO is again used -- albeit operating in a slightly different mode from that used for... ..difference, as will be seen, is due to the fact that pairs of the sub-processor systems 10 can be operated in a synchronized, lock-step mode, called duplex mode, in which each CPU 12 operates to execute the... ..1A illustrates another feature of the invention: a cross-link connection between the two sub-processor systems 10A, 10B through the use of additional routers 14 (identified in Fig. 1A as... ..RY(sub(1)), and RY(sub(2)) form a cross-link connection between the sub-processors 10A, 10B (or,

as shown, "sides" X and Y, respectively) to couple them to I... ..the routers RX(sub(2)) and RY(sub(2)) provide the I/O packet interface units 16x and 16y with a dual ported interface. Of course, it will now be evident... ..lend themselves to being used in a manner that can extend the configuration of the processing system 10 to include additional sub-processor systems such as illustrated in Figs. 1B and 1C. In Fig. 1B, for example, one... ..of each of the routers 14A and 14B is used to connect the corresponding sub-processor systems 10A and 10B to additional sub-processor systems 10A' and 10B' forming thereby a larger processing system comprising clusters of the basic processing system 10 of Fig. 1.

Similarly, in Fig. 1C the above concept is extended to form an eight sub-processor system cluster, comprising sub-processor systems pairs 10A/10B, 10A'/10B', 10A"/10B", and 10A'''/10B'''. In turn, each of the sub-processor systems (e.g., sub-processor system 10A) will have essentially the same basic minimum configuration of a CPU 12, a... ..by a I/O packet interface 16, except that, as Fig. 1C shows, the sub-processor systems 10A and 10B include additional routers 14C and 14D, respectively, in order to extend the cluster beyond sub-processor systems 10A'/10B' to the sub-processor systems 10A"/10B" and 10A'''/10B'''. As Fig. 1C further illustrates, unused ports 4 and... ..the routers 14 when configuring the topology of the system 10, any CPU 12 of processing system 10 of Fig. 1C can access any other "end unit" (e.g., a CPU or I/O device) of any of the other sub-processor systems. Two paths are available from any CPU 12 to the last router 14 connecting to the I/O packet interface 16. For example, the CPU 12B of the sub-processor system 10B' can access the I/O 16''' of sub-processor system 10A''' via router 14B (of sub-processor system 10B'), router 14D, and router 14B (of sub-system 10B'') and, via link LA... ..10A'''), OR via router 14A (of sub-system 10A'), router 14C, and router 14A (sub-processor system 10A'''). Similarly, CPU 12A of sub-processor system 10A" may access (via two paths) memory contained in the CPU 12B of sub-processor 10B to read or write data. (Memory accesses by one CPU 12 of another component of the processing system requires, as will be seen, the components seeking access to have authorization to do... ..prevents corruption of memory data of a CPU by erroneous access.)

The topology of the processing system shown in Fig. 1B is achieved by using port 1 of the routers 14A, 14B, and auxiliary TNet links LA, to connect to the routers 14A', 14B' of sub-processor systems 10A', 10B'. The topology thereby obtained establishes redundant communication paths between any CPU 12 (12A, 12B, 12A', 12B') and any I/O packet interface 16 of the processing system 10 shown in Fig. 1B. For example, the CPU 12A' of the sub-processor system 10A' may access the I/O 16A of sub-processor system 10A by a first path formed by the router 14A' (in port 4, out... ..shown in Fig. 1B. By interconnecting one port of each router 14 of each sub-processor pair, and using additional auxiliary TNet links LA (illustrated in Fig. 1C with the dotted line connections) between the ports 1 of the routers 14 (14A" and 14B'') of sub-processor systems 10A", 10B" and 10A''', 10B''', two separate, independent data paths can be found between any CPU 12 and any I/O packet interface 16. In this fashion, any end unit (i.e., a CPU 12 or an I/O packet interface 16) will have at least two paths to any other end unit.

Providing alternate paths of access between any two end units (e.g., between a CPU 12 and any other CPU 12, or between any CPU... ..any two of the remaining fault domains. Here, a fault domain could be a sub-processor system (e.g., 10A). Thus, if the sub-processor system 10A were brought down because of a failure the electrical power being supplied, without... ..TNet link LA between the routers 14A''' and 14B''', the CPU 12B of the sub-processor system 10B would have lost access to the I/O packet interface 16''' (via router... ..with the loss of the router 14A (and router 14C) by loss of the sub-processor system 10A, communications between the CPU 12B is still possible via the route of router... ..equally to CPU 12B. As Fig. 2 shows, the CPU 12A includes a pair of processor units 20a, 20b that are configured for synchronized, lock-step operation in that both processor units 20a, 20b receive and execute identical instructions, and issue identical data and command outputs, at substantially the same moments in time. Each of the processor units 20a and 20b is connected, by a bus 21 (21a, 21b) to a corresponding cache memory 22. The particular type of processor units used could contain sufficient internal cache memory so that the cache memory 22 would not... ..22 could be used to supplement any cache memory that may be internal to the processor units 20. In any event, if the cache memory 22 is used, the bus 21 is... ..22 address bits, 3 bits of parity covering the address, and 7 control bits.

The processors 20a, 20b are also respectively coupled, via a separate 64-bit address/data bus 23 to X and Y interface units 24a, 24b. If desired, the address/data communicated on each bus 23a, 23b could also be protected by parity, although this will increase the width of the bus. (Preferably, the processors 20 are constructed to include RISC R4000 type microprocessors, such as are available from the MIPS Division of Silicon Graphics, Inc. of Santa Clara, California.)

The X and Y interface units 24a, 24b operate to communicate data and command signals between the processor units 20a, 20b and a memory system of the CPU 12A, comprising a memory controller (MC... ..MC halves 26a and 26b) and a dynamic random access memory array 28. The interface units 24 interconnect to each other and to the MCs 26a, 26b by a 72-bit... ..accompanied by 8 bits of ECC) are written to the memory 28 by the interface units 24, one interface unit 24 will drive only one word (e.g., the 32 most significant portion) of the doubleword being written while the other interface unit 24 writes the other word of the double word (e.g., the least significant 32-bit portion of the doubleword). In addition, on each write operation the interface units 24a, 24b perform a cross-check operation on the data not written by that interface unit 24 with the data written by the other to check for errors; on read operations... ..accessed corresponds to the address of the location from which the doubleword was stored.

Interface units 24a, 24b of the CPU 12A form the circuitry to respectively service the X and Y (I/O) ports of the CPU 12A. Thus, the X interface unit 24a connects by the bi-directional TNet Link Lx to a port of the router 14A of the processor system 10A (Fig. 1A) while the Y interface unit 24b similarly connects to the router 14B of the processor system 10B by TNet Link Ly. The X interface unit 24a handles all I/O traffic between the router 14A and the CPU 12A of the sub-processor system 10A. Likewise, the Y interface unit 24b is responsible for all I/O traffic between the CPU 12A and the router 14B of companion sub-processor system 10B.

The TNet Link Lx connecting the X interface unit 24a to the router 14A (Fig. 1) comprises, as above indicated, two 10-bit buses... ..sub(x)) carries data incoming from the router 14A. In similar fashion, the Y interface unit 24b is connected to the router 14B (of the sub-processor system 10B) by two 10-bit buses: 30(sub(y)) (for outgoing transmissions) and 32... ..y)) (for incoming transmissions), together forming the TNet Link Ly.

The X and Y interface units 24a, 24b are synchronously operated in lock-step, performing substantially the same operations at substantially the same times. Thus, although only the X interface unit 24a actually transmits data onto the bus 30(sub(x)), the same output data is being produced by the Y interface unit 24b, and used for error-checking. The Y interface unit 24b output data is coupled to the X interface unit 24a by a cross-link 34(sub(y)) where it is received by the X interface unit 24a and compared against the same output data produced by the X interface unit. In this way the outgoing data made available at the X port of the CPU... ..the port of the CPU 12A is checked. The output data from the Y interface unit 24b is coupled to the Y port by a 10-bit bus 30(sub(y)), and also to the X interface unit 24a by the 9-bit cross-link 34(sub(y)) where it is checked with that produced by the X interface unit.

As mentioned, the two interface units 24a, 24b operate in synchronous, lock-step with one another, each performing substantially the same... ..X and/or Y ports of the CPU 12A must be received by both interface units 24a, 24b to maintain the two interface units in this lock-step mode. Thus, data received by one interface unit 24a, 24b is passed to the other, as indicated by the dotted lines and 9... ..sub(x)) (communicating incoming data being received at the X port by the X interface unit 24a to the Y interface unit 24b) and 36(sub(y)) (communicating data received at the Y port by the Y interface unit 24b to the X interface unit 24a).

Certain more robust operating systems are structured with a fault-tolerant capability in the... ..example, U.S. Patent No. 4,817,091 teaches a multiprocessor system in which each processor periodically messages each of the processors of the system (including itself), under software control, to thereby provide an indication of continuing operation. Each of the processors, in addition to performing its normal tasks, operates as a backup processor to another of the processors. In the event one of the backup processors fails to receive the messaged indication from a sibling processor, it will take over the operation of that sibling (now thought to be inoperative), in... ..platform for

both types of software. Thus, when a robust operating system is available, the processing system 10 can be configured to operate in a "simplex" mode in which each of... ..left, in most instances, to software.

Alternatively, for less robust operating systems and software, the processing system 10 provides a hardware-based fault-tolerance by being configured to operate in a... ..g., CPUs 12A, 12B) are coupled together as shown in Fig. 1A, to operate in synchronized, lock-step fashion, executing the same instructions at the substantially the same moment in time... ..data and command symbols. In order to simplify the design of the CPU 12, the processors 20 are precluded from communicating directly with any outside entity (e.g., another CPU 12... ..O device via the I/O packet interface 16). Rather, as will be seen, the processor will construct a data structure in memory and turn over control to the interface units 24. Each interface unit 24 includes a block transfer engine (BTE; Fig. 5) configured to provide a form of... ..to the destination according to information contained in the message packet.

The design of the processing system 10 permits a memory 28 of a CPU to be read or written by... ..via the routers 14. Accordingly, before continuing with the description of the construction of the processing system 10, it would be of advantage to understand first the configuration of the data...information.

As indicated, the HADC message packet operates to communicate write data between the end units (e.g., CPU 12) of the processing system 10. Other message packets, however, may be differently constructed because of their function and... ..CRC. The HC message packet is used to acknowledge a request to write data.

Interface Unit:

The X and Y interface units 24 (i.e., 24a and 24b - Fig. 2) operate to perform three major functions within the CPU 12: to interface the processors 20 to the memory 28; to provide an I/O service that operates transparently to, but under the control of, the processors; and to validate requests for access to the memory 28 from outside sources.

Regarding first the interface function, the X and Y interface units 24a, 24b operate to respectively communicate processors 20a, 20b to the memory controllers (Mcs 26a, 26b) and memory 28 for writing and... ..fast checking of the data read/written. For example, write operations have the two interface units 24a, 24b cooperating to cross-check the data to be written to ensure its integrity (and at the same time, the interface units 24 will operate) to develop an error correcting code (ECC) that covers, as will be With respect to I/O access, the processors 20 are not provided with the ability to communicate directly with the input/output systems... ..must write data structures to the memory 28 and then pass control to the interface units 24 which perform a direct memory access (DMA) operation to retrieve those data structures, and... ..indicated in the data structure itself.)

The third function of the X and Y interface units 24, access validation to the memory 28, uses an address validation and translation (AVT) table maintained by the interface units. The AVT table contains an address for each system component (e.g., an I/O... ..the incoming message packets are virtual addresses. These virtual addresses are translated by the interface unit to physical addresses recognizable by the memory control units 26 for accessing the memory 28.

Referring to Fig. 5, illustrated is a simplified block diagram of the X interface unit 24a of the CPU 12A. The companion Y interface unit 24b (as well as the interface units 24 of the CPU 12B, or any other CPU 12) is of substantially identical construction. Accordingly, it will be understood that a description of the interface unit 24a will apply equally to the other interface units 24 of the processing system 10.

As Fig. 5 illustrates, the X interface unit 24a includes a processor interface 60, a memory interface 70, interrupt logic 86, a block transfer engine (BTE) 88, access validation and translation logic 90, a packet transmitter 94, and a packet receiver 96.

Processor Interface:

The processor interface 60 handles the information flow (data and commands) between the processor 20a and the X interface unit 24a. A processor bus 23, including a 64 bit address and data bus (SysAD) 23a and a 9 bit command

bus 23b, couples the processor 20a and the processor interface 60 to one another. While the SysAD bus 23a carries memory address and data... ..and qualifying commands carried at substantially the same time on the SysAD bus 23a. The processor interface 60 operates to interpret commands issued by the processor unit 20a in order to pass reads/writes to memory or control registers of the processor interface. In addition, the processor interface 60 contains temporary storage (not shown) for buffering addresses and data for access to... ..26). Data and command information read from memory is similarly buffered en route to the processor unit 20a, and made available when the processor unit is ready to accept it. Further, the processor interface 60 will operate to generate the necessary interrupt signalling for the X interface unit 24a.

The processor interface 60 is connected to a memory interface 70 and to configuration registers 74 by a bi-directional 64 bit processor address/data bus 76. The configuration registers 74 are a symbolic representation of the various control registers contained in other components of the X interface unit 24a, and will be discussed when those particular components are discussed. However, although not specifically... ..throughout other of the logic that is used to implement the X interface 24a, the

processor address/data bus 76 is likewise coupled to read or write to those registers.

Configuration registers 74 are read/write accessible to the processor 20a; they allow the X interface unit to be "personalized." For example, one register identifies the node address of the CPU 12A... ..with the CPU 12A; another, readable only, contains a fixed identification number of the interface unit 24, and still other registers define areas of memory that can be used by, for... ..logic 90, etc. employing them are discussed.

The memory interface 70 couples the X interface unit 24a to the memory controllers 26 (and to the Y interface unit 24b; see fig. 2) by a bus 25 that includes two 36 bi-directional bit... ..25a, 25b. The memory interface operates to arbitrate between requests for memory access from the processor unit 20, the BTE 88, and the AVT logic 90. In addition to memory accesses from the processor unit 20a, the memory 28 may also be accessed by components of the processing system 10 to, for example, store data requested to be read by the processor unit 20a from an I/O unit 17, or memory 28 may also be accessed for I/O data structures previously set up in memory by the processor unit. Since these accesses are all asynchronous, they must be arbitrated, and the memory interface 70... ..command information accessed from the memory 28 is coupled from the memory interface to the processor interface 60 by a memory read bus 82, as well as to an interrupt logic... ..doubleword quantities. However, while the memory interfaces 70 of both the X and Y interface units 24a and 24b ...by the memory interface 70 are coupled to the memory interface by the companion interface unit 24 where they are compared with the same 32 bits for error.

Digressing for the... ..containing interrupt information are received, that information is conveyed to the interrupt logic 86 for processing and posting for action by the processor 20, along with any interrupts generated internal to the CPU 12A. Internally generated interrupts will... ..register 71 (internal to the interrupt logic 86), indicating the cause of the interrupt. The processor 20 can then read and act upon the interrupt. The interrupt logic is discussed more fully below.

The BTE 88 of the X interface unit 24a operates to perform direct memory accesses, and provides the mechanism that allows the processors 20 to access external resources. The BTE 88 can be set-up by the processors 20 to generate I/O requests, transparent to the processors 20 and notify the processors when the requests are complete. The BTE logic 88 is discussed further below.

Requests for... ..8 byte wide format necessary for storing in the memory 28.

Outgoing message packets containing processor originated transaction requests (e.g., a read request asking for a block data from an I/O unit) are monitored by the request transaction logic (RTL) 100. The RTL 100 provides a time... ..will generate an interrupt (handled and reported by the interrupt logic 86) to inform the processor 20 that

the request was not honored. In addition, the RTL 100 will validate responses... 28 (by the DMA operation of the BTE 86) at a location known to the processor 20 so that it can locate the response.

Each of the CPUs 12 are checked... discussed. One such check is an on-going monitor of the operation of the interface units 24a, 24b of each CPU. Since the interface units 24a, 24b operate in lock-step synchronism checking can be performed by monitoring the operating states of the paired interface units 24a, 24b by a continuous comparison of certain of their internal states. This approach is implemented by using one stage of a state machine (not shown) contained in the unit 24a of CPU 12A, and comparing each state assumed by that stage with its identical state machine stage in the interface unit 24b. All units of the interface units 24 use state machines to control their operations. Preferably, therefore, a state machine of the memory interface 70 that controls the data transfers between the interface unit 24 and the MC 26 is used. Thus, a selected stage of the state machine used in the memory interface 70 of the interface unit 24a is selected. An identical stage of a state machine of one of the interface unit 24b is also selected. The two selected stages are communicated between the interface units 24a, 24b and received by a compare circuit contained in both interface units 24a, 24b. As the interface units operate lock-step with one another, the state machines will likewise march through the same identical states, assuming each state at substantially the same moments in time. If an interface unit encounters an error, or fails, that activity will cause the interface units to diverge, and the state machines will assume different states. The time will come when... that will bring to the attention of the CPUs 12A (or 12B) that the interface units 24a, 24b of that CPU are no longer in lock-step, and to act accordingly... X port, receiving only those message packets transmitted by the router 14A of the sub-processor system 10A (Fig. 1A). The Y port is serviced by the Y interface unit 24b to receive message packets from the router 14B of the companion sub-processor system 10B. However, both interfaces (as well as Mcs 26 and processor 20), as has been indicated, are basically mirror images of one another in that both... in both structure and function. For this reason, message packet information, received by one interface unit (e.g., 24a) must be passed for processing also to the companion interface unit (e.g., 24b). Further, since both interface units 24a, 24b will assemble the same message packets for transmission from the X or the Y ports, the message packet being transmitted by the interface unit (e.g., 24b) actually being communicated from the associated port (e.g., the Y port) will also be coupled to the other interface unit (e.g., 24a) for cross-checking for errors. These features are illustrated in Figs. 6... receiving portions of the packet receivers 96 (96x, 96y) of the X and Y interface units 24a, 24b are broadly illustrated. As shown, each packet receiver 96x, 96y has a clock... receive a corresponding one of the TNet Links 32. The CS FIFOs 102 operate to synchronize the incoming command/data symbols to the local clock of the packet receiver 96, buffering... 104x, coupled to the MUX 104y of the packet receiver 96y of the Y interface unit 24b by the cross-link connection 36(sub(x)). In similar fashion, information received at the Y port is coupled to the X interface unit 24a by the cross-link connection 36(sub(y)). In this manner, the command/data... packets received at one of the X, Y ports by the corresponding X, Y, interface unit 24a, 24b is passed to the other so that both will process and communicate the same information on to other components of the interface units 24 and/or memory 28.

Continuing with Fig. 6, depending upon which port X, Y... or the other of the CS FIFOs 102x, 102y for communication to the storage and processing logic 110 of the interface unit 24. The information contained in each 9-bit symbol is an 8-bit byte of... the encoding of which is discussed below with respect to Fig. 9. The storage and processing logic 110 will first translate the 9-bit symbols to 8-bit data or command... the outputs of the CS FIFOs 102x, 102y are also coupled to a command decode unit in addition to the MUX 104. The command decode unit operates to recognize command symbols (differentiating them from data symbols in a manner that is... below), decoding them to generate therefrom command signals that are applied to a receiver control unit, a state machine-based element that functions to control packet receiver operations.

As indicated above... at the output of the MUX 104, the receiver control portion of the storage control unit enables CRC check logic 106 to calculate a CRC symbol while the data symbols are... below, CS FIFOs are found not only in the packet receivers 96 of the interface units 24, but also at each receiving port of the routers 14 and the I/O... an even more important part, and perform a unique function, when a pair of sub-processor systems are operating in duplex mode and the two CPUs 12A and 12B of the sub-processor systems 10A, 10B operate in synchronized,

lock-step, executing the same instructions at the same time. When operating in this latter... ..difficult to ensure that the clocking regime of the routers 14A and 14B are exactly synchronized to those of the CPUs 12A and 12B - even when using frequency locked clocking. In... ..used to transmit symbols to a CPU 12 and the clock used by an interface unit 24 to receive those symbols.

The structure of the CS FIFO 102 is diagrammatically illustrated... ..i.e., a packet) or IDLE symbols - except during certain situations (e.g., reset, initialization, synchronization and others discussed below). As explained above, each symbol held in the transmit register 120...same symbol leaving the storage queue, allowing each symbol entering the storage queue 126 to settle before it is clocked out and passed to the storage and processing units 110x (and 110y) by the MUX 104x (and 104y). Since the transmit and receive clocks... ..functioning in duplex mode) operate to transmit symbols with near frequency clocking. Even so, clock synchronization FIFOs are used at these other ports to receive symbols transmitted with near frequency clocking, and the structure of these clock synchronization FIFOs are substantially the same as that used in frequency locked environments, i.e., that... ..of the storage queue 126 are nine bits wide; in near frequency environments, the clock synchronization FIFOs use symbol locations of the queue 126 that are 10 bits wide, the extra... ..the faster clock source. To handle this clock drift, the two pointers are effectively re-synchronized periodically.

When the CPUs 12 are paired and operating in duplex mode, all four interface

units 24 operate in lock-step to, among other things, transmit the same data and receive... ..simplex mode, each independent of the other, clocking need only be near frequency.

The interface unit 24 receives a SYNC CLK signal that is used in combination with a SYNC command symbol to initialize and synchronize the Rcv register 124 to the transmitting router 14. When using either near frequency or... ..102X preferably begin from some known state. Incoming symbols are examined by the storage and processing units 110 of the packet receivers 96. The storage and processing units look for, and act upon as appropriate, command symbols. Pertinent here is that when the... ..receives a SYNC command symbol it will be decoded and detected by the storage and processing unit 110. Detection of the SYNC command symbol by the storage and processing unit 110 causes assertion of a RESET signal. The RESET signal, under synchronous control of the SYNC CLK signal, is used to reset the input buffers (including the clock synchronization buffers) to predetermined states, and synchronize them to the routers 14.

The synchronization of the CS FIFOs 102 of the interface units 24 those of one or both routers 14A, 14B is discussed more fully below in the section discussing synchronization.

Packet Transmitter:

Each interface unit 24 is assigned to transmit from and receive at only one of the X or Y ports of the CPU 12. When one of the interface units 24 transmits, the other operates to check the data being transmitted. This is an important... ..shows, in abbreviated form, the packet transmitters 94x, 94y of the X and Y interface units 24a, 24b, respectively. Both packet transmitters are identically constructed, so that discussion of one (packet... ..logic 152 that receives, from the BTE 88 or AVT 90 of the associated interface unit (here, the X interface unit 24a) the data to be transmitted - in doubleword (64-bit) format. The packet assembly logic... ..and Y ports: they are either symbols that make up a message packet in the process of being transmitted, or IDLE symbols, or other command symbols used to perform control functions... ..154, 156. The output of the multiplexer 154 connects to the X port. (The interface unit 24b connects the output of the multiplexer 154 to the Y port.) The multiplexer 156... ..sub(x)) to the checker logic 160 of the packet transmitter 94y (of the interface unit 24b).

A selection (S) input of the multiplexers receives a 1-bit output from an... ..is accessible to the MP 18 via an OLAP (not shown) formed in the interface unit 24, and is written with information that "personalizes," among other things, the interface units 24 Here, the X/Y stage of the configuration register 162 configures the packet transmitter 94x of

the X interface unit 24a to communicate the X encoder 150x output to the X port; the output of... traffic is present, the operation of the two packet interfaces 94 (and, thereby, the interface units 24 with which they are associated) are continually monitored. Should one of the checkers detect... will be asserted, resulting in an internal interrupt being posted for appropriate action by the processors 20.

Message packet traffic operates in the same manner. Assume, for the moment, that the... that information, a byte at a time, to the X encoder 150x of both interface units 96, which will translate each byte to encoded 9-bit form. The output of the... is checked with that from the packet transmitter 94x. Again, the operation of the interface units 24a, 24b, and the packet transmitters they contain, are inspected for error.

In the same... monitored.

Returning for the moment to Fig. 5, if the outgoing message packet is a processor initiated transaction (e.g., a read request), the processors 20 will expect a message packet to be returned in response. Thus, when the BTE... will issue a timeout signal to the interrupt logic (Fig. 14A) to thereby notify the processors 20 of the absence of a response to a particular transaction (e.g., a read... the access, to name just a few. Also, the area of memory of the memory unit 28 desired to be accessed are identified in the message packets by virtual or I... virtual addresses be translated to physical addresses of the memory 28. Finally, interrupts generated by units or elements external to the CPU 12A, are transmitted via message packets to interrupt the processors 20, which are also written to memory 28 when received. All this is handled by the interrupt logic and AVT logic 86, 90.

The AVT logic unit 90 utilizes a table (maintained by the processor 20 in memory 28) containing AVT entries for each possible external source permitted access to the memory 28. Each AVT entry identifies a specific source element or unit and the particular page (a page being nominally 4K (4096) bytes), or portion of a... expected" memory accesses. Expected memory accesses are those initiated by the CPU 12 (i.e., processors 20) such as a read request for information from an I/O device. These latter memory accesses are handled by a transaction sequence number (TSN) assigned to each processor initiated request. At about the time the read request is generated, the processors 20 will allocate an area of memory for the data expected to be received in... and 26b are, in turn, respectively coupled to the memory interfaces 70 of each interface unit 24a, 24b. The 64-bit doublewords are written to the memory 28 with the upper... check bits respectively from the memory interfaces 70 (70a, 70b) of each of the interface units 24a, 24b (Fig. 5).

Referring to Fig. 10, each memory interface 70 receives, from either the bus 82 from the processor interface 60 or the bus 83 from AVT logic 90 (see Fig. 5), of the associated interface unit 24, 64 bits of data to be written to memory. The busses 76 and 83... other for cross-checking between them. Thus, for example, the memory interface 70a (of interface unit 24a) will drive the MC 26a with the "upper" 32 bits of the 64 bits... are check bits, leaving 40 bits unused.

Access Validation:

As previously indicated, components of the processing system 10 external to the CPU 12A (e.g., devices of the I/O packet... not without qualification. Access validation, as implemented by the AVT logic 90 of the interface units 24, operates to prevent the content of the memory 28 from being corrupted by... Accesses to the memory 28 are validated by the AVT logic 90 of each interface unit 24 (Fig. 5), using all of six checks: (1) that the CRC of the message... also are permitted the particular message packet source.

The access validation mechanism of the interface unit 24a, AVT logic 88, is shown in greater detail in Fig. 11. Incoming message packets... and post an interrupt to the interrupt logic 86 (Fig. 5) for action by the processor 20.

The mask operation permits the size of the table of AVT entries to be varied. The content of the AVT mask register 175 is accessible to the processor 20, permitting the processors 20 to optionally select the size of the AVT entry table. A maximum AVT table... 172 allows the AVT size to be matched to the needs of the system. A processing system 10 that includes a larger number of external elements (e.g., the number of... amount of the memory space of

memory 28 to the AVT entries. Conversely, a smaller processing system 10, with a smaller number of external elements will not have such a large... ..set to a logic "ZERO" indicate an nonexistent TNet address, outside the limits of the processing system 10. A received packet with a TNet address outside the allowable TNet range will... ..in Fig. 11 as being held in the AVT entry register 180 during the validation process. AVT entries have two basic formats: normal and interrupt. The format of a normal AVT...of the AVT input register 170) will result in an error being posted to the processor via an interrupt.

A 12-bit "Permissions" field is included in t AVT entry to... ..path =0). Denials are logged as interrupts with the interrupt logic, and reported to the processor 20 - if the E field is set to a state ("ONE") that enables error-reporting... ..e.g., to a "ONE"), the other fields (Upper Bound, etc.) gain new definitions for processing interrupt writes and managing interrupt queues. This is discussed in more detail below in connection... ..memory 28 will be handled. Set to one state, the requested write operation will be processed normally; set to a second state, write requests specifying addresses with a fractional cache line...be written to a specific queue (interrupt queue) in memory 28, with signalling provided the processors 20 to indicate that an interrupt has been received and "posted," and ready for servicing by the processors 20. Since the interrupt queues are at specific memory locations, the processor can obtain the interrupt data when needed.

An AVT interrupt entry for an interrupt may... ..by the interrupt logic 86, and extracted from the head of the queue by the processor 20 when servicing the interrupt.

The AVT interrupt entry also includes a 20-bit segment ("Source ID") containing source ID information, identifying the external unit seeking attention by the interrupt process. If the source ID information of the AVT interrupt entry does not match that contained... ..class" of the interrupt that is used to determine the interrupt level set in the processor 20 (described more fully below); (2) a queue number that is used to select, as...capability to deliver interrupts to a CPU 12 for servicing. For example, an I/O unit may be unable to complete a read or write transaction issued by a CPU because... ..identify the recipient. These and other errors, exceptions, and irregularities, noted by the I/O units, or the I/O Interface elements, can become the a condition that requires the intervention... ..the AVT entry register 180 for use by the interrupt logic 86 of the interface unit 24 (Fig. 5), illustrated in greater detail in Fig. 14A.

It is interrupt logic 86...four circular queues specified by the base address information contained in the AVT entry. The processor(s) 20 will then be notified, and it will be up to them as to... ..selected tail queue register 256 by combiner circuit 270, the output of which is the processed by the "mod z" circuit 273 to turn new offset into the queue at which... ..signal. The Queue Full warning signal becomes an "intrinsic" interrupt that is conveyed to the processor units 20 as a warning that if the matter is not promptly handled, later-received interrupt... ..will be discarded.

Incoming message packet interrupts will cause interrupts to be posted to the processor 20 by first setting one of a number of bit positions of an interrupt register 280. Multi-entry queued interrupts are set in interrupt registers 280a for posting to the processor 20; single-entry queue interrupts use interrupt register 280b. Which bit is set depends upon... ..multi-entry queued interrupts, soon after a multi-entry queued interrupt is determined, the interface unit will assert a corresponding interrupt signal (I1) that is applied to decode circuit 283. Decode... ..of register 280a to set, thereby providing advance information concerning the received interrupt to the processor(s) 20, i.e., (1) the type of interrupt posted, and (2) the class of... ..to one another by a compare circuit 279. The update register is writable by the processor 20 to select a register pair for comparison. If the content of the two selected... ..cleared.

Digressing for the moment, there are two basic types of interrupts that concern the processors 20: those interrupts that are communicated to the CPU 12 by message packets, and those...the seven interrupt postings to a latch 288, from which they are coupled to the processor 20 (20a,20b) which has an interrupt register for receiving holding the postings.

In addition... ..change in interrupts (either an interrupt has been serviced, and its posting deleted by the processor 20, or a new interrupt has been posted), a "CHANGE" signal will be issued to the processor interface 60 to inform it that an interrupt posting change has occurred, and that it should communicate the change to the processor 20.

Preferably, the AVT entry register 180 is configured to operate like a single line... ..such as set-associative, fully-associate, or direct-mapped, to name a few.

Coherency:

Data processing systems that use cache memory have long recognized the problem of coherency: making sure that... ..the incoming packet is permitted access are applied to a boundary crossing (Bdry Xing) check unit 219. Boundary check unit 219 also receives an indication of the size of the cache block the CPU 12... ..Len field of the header information from the AVT input register 170. The Bdry Xing unit determines if the data of the incoming packet is not aligned on a cache boundary...time an interrupt will be written to the queued interrupt register 280, to alert the processors 20 that a portion of the incoming data is located in the special queue.

In... ..not, the packet (both header and data) is written to a special queue, and the processors so notified by the intrinsic interrupt process described above. The processors may then move the data from the special queue to cache 22, and later write... ..the cache 22 and the memory 28 is preserved.

Block Transfer Engine (BTE):

Since the processor 20 is inhibited from directly communicating (i.e., sending) information to elements external to the... ..indirect method of information transmission.

The BTE 88 is the mechanism used to implement all processor initiated I/O traffic to transfer blocks of information. The BTE 88 allows creation of... ..BTE registers 300, 302 whose content is coupled to the MUX 306 (of the interface unit 24a; Fig. 5) and used to access the system memory 28 via the memory controllers... ..BTE data structure 304 in the memory 28 of the CPU 12A (Fig. 2). The processors 20 will write a data structure 304 to the memory 28 each time information is... ..begin on a quadword boundary, and the BTE registers 300, 302 are writable by the processors 20 only. When a processor does write one of the BTE registers 300, 302, it does so with a word... ..the request bit (rc0, rc1) to a clear state, which operates to initiate the BTE process, which is controlled by the BTE state machine 307.

The BTE registers 300, 302 also... ..cause (ec) bit differentiates time-outs and NAKs.

When information is being transferred by the processors 20 to an external unit, the data buffer portion 304b of the data structure 304 holds the information to be transferred. When information from an external unit is received by the processors 20, the data buffer portion 304b is the location targeted to hold the read response information.

The beginning of the data structure 304, portion 304a written by the processor 20, includes an information field (Dest), identifying the external element which will receive the packet... ..the transmitted data is to be written. This information is used by the packet transmitter unit 120 (Fig. 5) to assemble the packet in the form shown in Figs. 3-4...list (e1) bit, when set, indicates the end of the chain, and halts the BTE processing.

The interrupt completion (ic) bit, when set, will cause the interface unit 24a to assert an interrupt (BTECmp) which sets a bit in the interrupt register 280... ..the chain pointer).

The interrupt time-out (it) bit, when set, will cause the interface unit 24a to assert an interrupt signal for the processor 20 if the acknowledgement of the access times-out (i.e., if the request timer... ..time), or elicits a NAK response (indicating that the target of the request could not process the request).

Finally, if the check sum (cs) bit is set, the data to be... ..containing the data from which the check sum was formed.

To sum up, when the processors 20 of the CPU 12A desire to send data to an external unit, they will write a data structure 304 to the memory 28, comprising identifier information in portion 304a of the data structure, and the data in the buffer portion 304b. The processors 20 will then determine the priority of the data and will write the BTE register... information, and sent.

If the data structure 304 indicates a read request (i.e., the processors 20 are seeking data from an external unit - either an I/O device or a CPU 12), the Len and Local Buffer Ptr... receiver 100 (Fig. 5) until the local memory write operation is executed.

Responses to a processor-generated read request to an external unit are not processed by the AVT table logic 146. Rather, when the processors 20 set up the BTE data structure, a transaction sequence number (TSN) is assigned the... the BTE 88, which will be an HAC type packet (Fig. 4) discussed above. The processors 20 will also include an memory address in the BTE data structure at which the...302, assume that the foregoing transfer of data from the CPU 12A to an external unit is of a large block of information. Accordingly, a number of data structures would be set up in memory 28 by the processors 20, each (except the last) including a chain pointer to additional data structures, the sum... sent. Assume now that a higher priority request is desired to be made by the processors 20. In such a case, the associated data structure 304 for such higher priority request... with another BTE operation descriptor.

Memory Controller:

Returning, for the moment, to Fig. 2, interface units 24a, 24b access the memory 28 via a pair of memory controllers (MC) 26a, 26b. The Mcs provide a fail-fast interface between the interface units 24 and the memory 28. The Mcs provide the control logic necessary for accessing... in dynamic random access memory (DRAM) logic). The Mcs receive memory requests from the interface units 24, and execute reads and writes as well as providing refresh signals to the DRAMs... to provide a 72 bit data path between the memory array 28 and the interface units 24a, 24b, which utilize an SBC-DBD-SbD ECC scheme, where b=4, on a... 26a, 26b to work together and simultaneously supply a 64-bit word to the interface units 24 with minimum latency, one-half of which (DO) comes from the MC 26a, and the other half (D1) comes from the other MC 26b. The interface unit 24 generate and check the ECC check bits. The ECC scheme used will not only... 26 bus 25, as well as in internal registers.

From the viewpoint of the interface units 24, the memory 28 is accessed with two instructions: a "read N doubleword" and a... doubleword read or a block read format. The signal called "data valid" tells the interface units 24 two cycles ahead of time that read data is being returned or not being returned.

As indicated above, the maintenance processor (MP 18; Fig. 1A) has two means of access to the CPUs 12. One is... 18 will write a register contained in the OLAP 285 with instructions that permit the processors 20 to build an image of a sequence of instructions in the memory that will permit them (the processors 20) to commence operation, ...to transfer instructions and data from an external (storage) device that will complete the boot process.

The OLAP 285 is also used by the processors 20 to communicate to the MP 18 error indications. For example, if one of the interface units 24 detect a parity error in data received from the memory controller 26, it will... and address transfers on the bus 25 between the MC 26a and the corresponding interface unit 24a. The addressing and data transfers on the DRAM data bus, as well as generation... the CPU 12.

Packet Routing:

The message packets communicated between the various elements of the processing system 10 (e.g., CPUs 12A, 12B, and devices coupled to the I/O packet... First, each TNet Link L connects to an element (e.g., router 14A) of the processing system 10 via a port that has both receive and transmit capability. Each transmit port... cycle (i.e., each clock period) of the T(underscore)Clk so that the clock

synchronization FIFO at the receiving end of the transmission will maintain synchronization.

Clock synchronization is dependent upon the mode in which the processing system 10 is operated. If operating in the simplex mode in which the CPUs 12A... connect directly to the CPUs may drift with respect to each other. Conversely, when the processing system 10 operates in a duplex mode (e.g., the CPUs operate in synchronized, lock-step operation), the clocks between routers 14 and the CPUs 12 to which they... not necessarily phase-locked).

The flow of data packets between the various elements of the processing system 10 is controlled by command symbols, which may appear at any time, even within... initiated by a CPU 12, or MP 18, and promulgated to all elements of the processing system 10 by the routers 14 to communicate an event requiring software action by all... command symbol is used in conjunction with near frequency operation as an aid to maintaining synchronization between the two clock signals that (1) transfer each symbol to, and load it in each receiving clock synchronization FIFO, and (2) that retrieves symbols from the FIFO.

SLEEP: This command symbol is sent by any element of the processing system 10 to indicate that no additional packet (after the one currently being transmitted, if... received.

SOFT RESET (SRST): The SRST command symbol is used as a trigger during the processes ("synchronization" and "reintegration," described below) that are used to synchronize symbol transfers between the CPUs 12 and the routers 14A, 14B, and then to place... SYNC command symbol is sent by a router 14 to the CPU 12 of the processing system 10 (i.e., the sub-processor systems 10A/10B) to establish frequency-lock synchronization between CPUs 12 and routers 14A, 14B prior to entering duplex mode, or when in duplex mode to request synchronization, as will be discussed more fully below. The SYNC command symbol is used in conjunction... or duplex to simplex), among other things, as discussed further below in the section on Synchronization and Reintegration.

THIS LINK BAD (TLB): When any system element receiving a symbol from a TNet link L (e.g., a router, a CPU, or an I/O unit) notes an error when receiving a command symbol or packet, it will send a TLB... identical pairs of symbols that are compared to one another when pulled from the clock synchronization FIFOs. The DVRG command symbol signals the CPU 12 that a mis-compare has been noted. When received by the CPUs, a divergence detection process is entered whereby a determination is made by the CPUs which CPU may be failing... command symbols described above operate to control message flow between the various elements of the processing system 10 (e.g., CPUs 12, router 14, and the like), using principally the BUSY end node" (i.e., a CPU 12 or I/O unit 17 - Fig. 1) may not assert backpressure because one of its transmit ports is backpressured... Improperly addressed packets are discarded by the router 14.

When a system element of the processing system 10 receives a BUSY command symbol on a TNet link L on which it... other command symbols (READY, BUSY, etc.).

Whenever a TNet port of an element of the processing system 10 detects receipt of a READY command symbol, it will terminate transmission of FILL... receives.

As will be seen, all elements (e.g., router 14, CPUs 12) of the processing system 10 that connect to a TNet link L for receiving transmitted symbols will receive those symbols via a clock synchronization (CS) FIFO. For example, as discussed above, the interface units 24 of CPUs 12 include all CS FIFOs 102x, 102y (illustrated in Fig. 6). The... depth to allow for speed matching, and the elastic FIFOs must provide sufficient depth for processing delays that may occur between transmission of a BUSY command symbol during receipt of a... another data byte in packet B. As packet A progresses to the next router, the process would be repeated. If the router 14 displaces more data bytes than the FIFO can... irrespective of its own findings.

SLEEP Protocol:

The SLEEP protocol is initiated by a maintenance processor via a maintenance interface (an on-line access port - OLAP), described below. The SLEEP protocol... ..reintegrate a slice of the system 10. Routers 14 must be idle (no packets in process) in order to change modes without causing data loss or corruption. When a SLEEP command symbol is received, the receiving element of processing system 10 inhibits initiation of transmission of any new packet on the associated transmit port... ..The HALT command symbol provides a mechanism for quickly informing all CPUs 12 in a processing system 10 that is necessary to terminate I/O activity (i.e., message transmissions between... ..CPUs that receive HALT command symbols on either of their receive ports (of the interface units 24) will post an interrupt to the interrupt register 280 if the system halt interrupt... ..interrupt; Fig. 14A).

The CPUs 12 may be provided with the ability to disable HALT processing. Thus, for example, the configuration registers 75 of the interface units 24 can include a "halt enable register" that, when set to a predetermined state (eg., ZERO) disables HALT processing, but reporting detection of a HALT symbol as an error.

Router Architecture:

Referring now to... ..simplified block diagram of the router 14A is illustrated. The other routers 14 of the processing system 10 (e.g., routers 14B, 14', etc.) are of substantially identical construction and, therefore...these ports 4, 5 are structured to operate in a frequency locked environment when a processing system 10 is set for duplex mode operation. In addition, when in duplex mode, a... ..5)) will receive the command/data symbols from the CPUs, pass them through the clock synchronization FIFOs 518 (discussed further below), and compare each symbol exiting the clock synchronization FIFOs with a gated compare circuit 517. When duplex operation is entered, a configuration register... ..517 to activate the symbol by symbol comparison of the symbols emanating from the two synchronization FIFOs 518 of the router input logic 502 for the ports 4 and 5. Of... ..to that received, at substantially the same time, by the other port input.

To maintain synchronization in the duplex mode, the two port outputs of the router 14A that transmit to... ..mode, are duplicated by the routers 14, and returned to both CPUs.) The output logic units 504(sub(4)), 504(sub(5)) that are coupled directly to the CPUs 12 will... ..message packet identifies only one of the duplexed CPUs 12, e.g., CPU 12A) in synchronized fashion, presenting those symbols in substantially simultaneous fashion to the two CPUs 12. Of course, the CPUs 12 (more accurately, the associated interface units 24) receive the transmitted symbols with synchronizing FIFOs of substantially the same structure as that illustrated in Fig. 7A so that, even... ..from the FIFO structures by both CPUs 12 on the same instruction cycle, maintaining the synchronized, lock-step operation of the CPUs 12 required by the duplex operating mode.

As will... ..conjunction with configuration data written to registers contained in control logic 509 by the maintenance processor 18 (via the on-line access port 285' and serial bus 19A; see Fig. 1A...links L. The input logic 505 of each port input 502 also assists in maintaining synchronization - at least for those ports sending symbols in the near-frequency environment - by removing received... ..slower-receiving element receiving symbols from a faster-sending element could overload the input clock synchronization FIFO of the slower-receiving element. That is, if a slower clock is used to pull symbols from the clock synchronization FIFO put there by a faster clock, ultimately the clock synchronization FIFO will overflow.

The preferred technique employed here is to periodically insert SKIP symbols in... ..stream to avoid, or at least minimize, the possibility of an overflow of the clock synchronization FIFO (i.e., clock synchronization FIFO 518; Fig. 20A) of a router 14 (or CPU 12) due to a T... ..being slightly higher in frequency than the local clock used to pull symbols from the synchronization FIFO. Using SKIP symbols to by-pass a push (onto the FIFO) operation has the... ..stall each time a SKIP command symbol is received so that, insofar as the clock synchronization FIFO is concerned, the transmitting clock that accompanied the SKIP symbol was missing.

Thus, logic... ..the port inputs 502 will recognize, and key off receipt of, SKIP command symbols for synchronization in the near frequency clocking environment so that nothing is pushed onto the FIFO, but... ..14, or between routers 14, or between a router 14 and an I/O interface unit 16A - Fig. 1) at a 50 Mhz rate, this allows for a

worst case frequency... symbol by supplying FILL or IDLE symbols (which are received and pushed onto the clock synchronization FIFOs, but are not passed to the elastic FIFOs). In short, each elastic FIFO 506...received symbols are then communicated from the input register 516 and applied to a clock synchronization FIFO 518, also by the T(underscore)Clk. The clock synchronization FIFO 518 is logically the same as that illustrated in Figs. 8A and 8B, used in the interface units 24 of the CPUs 12. Here, as Fig. 20A shows, the clock synchronization FIFO 518 comprises a plurality of registers 520 that receive, in parallel, the output of... 516. Associated with each of the registers 520 is a two-stage validity (V) bit synchronizer 522, shown in greater detail in Fig. 20B, and discussed below. The content of each... registers 520, together with the one-bit content of each associated two-stage validity bit synchronizer 522, are applied to a multiplexer 524, and the selected register/synchronizer pulled from the FIFO, and coupled to the elastic FIFO 506 by a pair of... is determined the state of the Push Select signal provided by a push pointer logic

unit 530; and, selection of which register 520 will supply its content, via the MUX 524... and loading of the register 520 selected by the push pointer logic 530. Similarly, the synchronization FIFO control logic 534 receives the clock signal local to the router (Rcv Clk) to... pointer logic 532.

Digressing for a moment, and referring to Fig. 20B, the validity bit synchronizer 522 is shown in greater detail as including a D-type flip-flop 541 with... 530 (Fig. 20A) selects the register 520 of the FIFO with which the validity bit synchronizer is associated for receipt of the next symbol - if not a SKIP symbol.

The delay... Truth Table, below). The D-type flip-flop 543 acts as an additional stage of synchronization, ensuring a stable level at the V output relative to the local Rec Clk. The... flip-flop 542, allowing the Pull signal (a periodic pulse from the sync FIFO Control unit 534) to clear the validity bit on this validity synchronizer 522 when the associated register 520 has been read. (Table omitted)

In summary, the validity synchronizer 522 operates to assert a "valid" (V) signal when a symbol is loaded in a...blocked from being routed out a particular port because another message is already in the process of being routed out that port. However, that other message in turn is also blocked...an incoming message packet bound for the CPUs will be replicated by the crossbar logic unit by routing the message packet to both port output 504(sub(4)) and 504(sub... P) identifies which of path (X or Y) should be used for accessing two sub-processing the device.

The routers 14 provide a capability of constructing a large, versatile routing network for, for example, massively parallel processing architectures. Routers are configured according to their location (i.e., level) in the network by...j)) and 509(sub(k)) are such that bits "def" are used in the algorithmic process, then bits "abc" of the Region ID are compared to the content of the Device... the route to default register 509(sub(f)) to the final stage of the selection process: check logic 602. Check logic 602 operates to check the status of the port output...a lower level router, and may be located in one or another of the sub-processing systems 10A, 10B. Whether a router is an upper level or lower level router depends... of CPUs 12 and I/O devices 16 to one another, forming a massively parallel processing (MPP) system. Other such MPP systems may exist, and it is those routers configured as... captured. As soon as the message packet's Destination ID is so captured, the selection process begins, proceeding to the development of a target port address that will be used to...an error that will be posted to the MP18 via the router's (or interface unit's) OLAP for action.

Digressing, it should be appreciated that these protocol rules observed by the routers 14 are also observed by the CPUs 12 (i.e., interface units 24) and I/O packet interfaces 17.

Finally, when the router 14A is in the... directly with the CPUs 12A, 12B, and duplex mode is used, a duplex operation logic unit 638 is utilized to coordinate the port output connected to one of the CPUs 12A... was able to write instructions to the OLAP 285 that would be executed by the processors 20 to build a small memory image and routine to permit the CPU 12 to... the clock generation circuit design. There will be one clock generator circuit in

each sub-processor system 10A/10B (Fig. 1) to maintain synchronism. Designated generally with the reference numeral 650... used by the various elements (e.g. CPU 12, routers 14, etc.) of the sub-processor system containing the clock circuit 650 (e.g., 10A).

The clock generator 654 is shown... The 50 Mhz clock signals produced by the counter 663 are distributed throughout the sub-processor system where needed.

Turning now to Fig. 25, there is illustrated the interconnection and use... the clock circuits 650 used to develop synchronous clock signals for a pair of sub-processor systems 10A, 10B (Fig. 1) for frequency locked operation. As illustrated in Fig. 25, the two CPUs 12A and 12B of the sub-processor systems 10A, 10B each have a clock circuit 650, shown in Fig. 25 as clock... 654B of both CPUs 12. A driver and signal line 667 interconnects the two sub-processor systems to deliver the M(underscore)CLK signal developed by the oscillator circuit 652A to the clock generator 654B of the sub-processor system 10B. For fault isolation, and to maintain signal quality, the M(underscore)CLK signal is delivered to the clock generator 654A of the sub-processor system 10A through a separate driver and a loopback connection 668. The reason for the... the cable (not shown) will establish the connection shown in Fig. 25 between the sub-processor systems 10A, 10B; connected another way, the connections will be similar, but the oscillator 652B... Fig. 25, the M(underscore)CLK signal produced by the oscillator circuit 652A of sub-processing system 10A is used by both sub-processing systems 10A, 10B as their respective SYNC CLK signals and the various other clock signals... produced by the clock generators 654A, 654B. Thereby, the clock signals of the paired sub-processing systems 10A, 10B are synchronized for the frequency locked operation necessary for duplex mode.

The VCXOs 662 of the clock... This allows both clock generators 654A, 654B to continue to provide to the two sub-processing systems 10A, 10B clock signals in the face of improper operation of the oscillator circuit 652A, although the sub-processor systems may no longer be frequency-locked.

The LOCK signals asserted by the phase comparators... LOCK signal signifies that the 50 Mhz signals produced by a clock generator 654 are synchronized, both in phase and in frequency, to the M(underscore)CLK signal. Thus, if either... signal that accompanies the symbol stream, and is used to push symbols onto the clock synchronizing FIFO of the receiving element (router 14, or CPU 12) is substantially identical in frequency... not phase, to that of the receiving element used to pull symbols from the clock synchronization FIFOs. For example, referring to Fig. 23, which illustrates symbols being sent from the router... clock (Local Clk). The former (Rev Clk) is used to push symbols onto the clock synchronization FIFOs 126 of each CPU, whereas the latter is used to pull symbols from the... much higher frequency clock signal. In such situations provision must be made to ensure that synchronization is maintained between the two CPUs as to symbols pulled from the clock synchronization FIFOs 126 of each.

Here, a constant ratio clocking mechanism is used to control operation of the two clock synchronization FIFOs 126, providing the clock signal that pulls symbols from the two FIFOs at the... control mechanism is shown, designated with the reference numeral 70. As Fig. 26A illustrates, clock synchronization FIFO control mechanism 700 includes an pre-settable, multi-stage serial shift register 702, the... ratio of the clock signal at which symbols are communicated and pushed onto the clock synchronization FIFOs 126 to the frequency of the clock signal used locally. Here, a 15 stages... that will be used as the Local Clk signal to pull symbols from the clock synchronization FIFOs 126, and to operate (update) the pull pointer counter 130. The selected output is... of the CPU 12 to the clock signal used to push symbols onto the clock synchronization FIFO 126, Rev Clk, the serial shift register is preset so that M stages of... duplexed CPUs 12 with a 50 Mhz clock. Thus, symbols are pushed onto the clock synchronization FIFOs 126 of the CPUs at a 50 Mhz rate. Assume further that the clock... of the MUX 704, which produces the clock signal that pulls symbols from the clock synchronization FIFOs 126, Rev Clk, will contain, for each 100 ns period, five clock pulses. Thus... five symbols will be pushed onto, and five symbols will be pulled from, the clock synchronization FIFOs 126.

This example is symbolically shown in Fig. 26B, while the timing diagram shown... ..labelled "IN" in Fig. 27) of the Rev Clk will push symbols onto the clock synchronization FIFOs 126. During that same 100 ns period, the serial shift register 702 circulates a... ..clocks which would require additional storage (i.e., an increase in the size of the synchronization FIFO) and impose more latency.

The constant ratio clock circuit presented here (Figs. 26) is... ..frequency to a clock regime of a different, higher frequency. The use of a clock synchronization FIFO is necessary here for compensating effects of signal delays when operating in synchronized, duplexed mode to receive pairs of identical command/data symbols from two different sources. However...so long as there are at least two registers in the place of the clock synchronization FIFO. Transferring data from a higher-frequency clock regime to a lower frequency clock regime... ..a wide range of possible clock ratios.

I/O Packet Interface:

Each of the sub-processor systems 10A, 10B, etc. will have some input/output capability, implemented with various peripheral units, although it is conceivable that the I/O of other sub-processor systems would be available so that a sub-processing system may not necessarily have local I/O. In any event, if local I/O... ..device (e.g., a signal line) would be received by the I/O packet interface unit 16 and used to form an interrupt packet that is sent to the CPU 12... ..OLAP bus, configuration information.

On-Line Access Port:

The MP 18 connects to the interface unit 24, memory controller (MC) 26, routers 14, and I/O packet interfaces with interface signals... ..OLAP 258 is essentially the same, regardless of what element (e.g. router 14, interface unit 24, etc.) it is used with. Fig. 28 diagrammatically illustrates the general structure of the... ..circuit chip used to implement certain of the elements discussed herein. For example, each interface

unit 24, memory controller 26, and router 14 is implemented by an application specific integrated circuit... ..of the OLAP 158 shown in Fig. 28 describes the OLAP associated with the interface unit 24, the MC 26, and the router 14 of the system.

As Fig. 28 shows...asymmetric variables, a "soft-vote" (SV) logic element 900 (Fig. 30A) is provided each interface unit 24 of each CPU 12. As Fig. 30 illustrates, the SV logic elements 900 of each interface unit 24 are connected to one another by a 2-bit SV bus 902, comprising bus lines 902a and 902b. Bus lines 902a carry one-bit values from the interface units 24 of CPU 12A to those of CPU 12B. Conversely, bus line 902b carries one... ..the CPU 12A.

Illustrated in Fig. 30B, is the SV logic element 900a of interface unit 24a of CPU 12A. Each SV logic element 900 is substantially identical in construction and... ..900a should be understood as applying equally to the other logic elements 900a (of interface unit 24b, CPU 12A), and 900b (of the interface units 24a, 24b of CPU 12B) unless noted otherwise. As Fig. 30B illustrates, the SV logic... ..the logic elements 900a (as well as its own). In this manner the two interface units 24a, 24b of the CPU 12A can communicate asymmetrical variables to each other.

In a... ..to the remote register 907 of logic element 902a (and that of the other interface unit 24b).

The logic elements 902 form a part of the configuration registers 74 (Fig. 5). Thus, they may be written by the processor unit(s) 20 by communicating the necessary data/address information over at least a portion of... ..local and remote registers 906 and 907.

The MUX 914 operates to provide each interface unit 24 of CPU 12A with selective use of the bus line 902a for the SV logic elements 900a, or for communicating a BUS ERROR signal if encountered during the reintegration process (described below) used to bring a pair of CPUs 12 into lock-step, duplex operation... ..same time, write the enable registers 912 of the logic element 900 of both interface units 24 of each CPU. One of the two logic elements

900 of each CPU will... it is the output enable registers 912 associated with the logic elements 900 of interface units 24a of both CPUs 12A, 12B that are written to enable the associated drivers 916. Thus, the output registers 904 of the interface units 24a of each CPU will be communicated to the bus lines 902; that is, the... to the bus line 902a, while the output register associated with logic element 900b, interface unit 24a of CPU 12B is communicated to bus line 902b. The CPUs 12 will both... again written by each CPU, followed again by reading the remote input registers 907. This process is repeated, one bit at a time, until the entire variable is communicated from the... each CPU 12 to the remote input register of the other. Note that both interface units 24 of CPU 12B will receive the bit of asymmetric information.

One example of use... elements 900 are also used to communicate bus errors that may occur during the reintegration process to be described. When reintegration is being conducted, a REINT signal will be asserted. As... ERROR signal is selected by the MUX 914 and communicated to the bus line 902a.

Synchronization:

Proper operation of the sub-processing systems 10A, 10B (Figs. 1A, 2) whether operating independently (simplex mode), or paired and operating in synchronized lock-step (duplex mode), requires assurance that data communicated between the CPUs 12A, 12B and... routers 14A, 14B will be received properly, and that any initial content of the clock synchronization FIFOs 102 (of CPUs 12A, 12B; Fig. 5) and 519 (of routers 14A, 14B; Fig... erroneously interpreted as data or commands. The push and pull pointers of the various clock synchronization FIFOs 102 (in the CPUs 12) and 518 (in the routers 14) need to be... apart, and presetting the associated FIFO queues to some known state. This done, all clock synchronization FIFOs are initialized for near frequency operation. ...in order to properly implement the lock-step operation of duplex mode operation, the clock synchronization FIFOs must be synchronized to operate with the particular source from which they receive data in order accommodate any... 14A, 14B to the CPUs 12A, 12B must be accounted for. It is the clock synchronization FIFOs 102 of the paired CPUs 12 that operate to receive message packet symbols, adjust... and present symbols to the two CPUs in a simultaneous manner to maintain lock-step synchronization necessary for duplex mode operation.

In similar fashion, each symbol received by the routers 14A... the CPUs (which is discussed further hereinafter). Again, it is the function of the clock synchronization FIFOs 518 of the routers 14A, 14B that receive message packets from the CPUs 12... so that the symbols received from the two CPUs 12 are retrieved from the clock synchronization FIFOs simultaneously.

Before discussing how the clock synchronization FIFOs of the CPUs and routers are reset, initialized, and synchronized, an understanding of their operation to maintain synchronous lock-step duplex mode operation is believed helpful. Thus, referring for the moment to Fig. 23, the clock synchronization FIFOs 102 of the CPUs 12A, 12B that receive data, for example, from the router... underscore)Clk, from the router 14A to the CPU 12B.

Consider operation of the clock synchronization FIFOs 102(sub(x)), 102(sub(y)), to receive identical symbol streams during duplex operation... held by the push and pull pointer counters 128, 130 for the CPU 12A (interface unit 24a), and the content of each of the four storage locations (byte 0, ..., byte 3... 6 show the same thing for the FIFO 102(sub(y)) of CPU 12B interface unit 24a for each symbol of the duplicated symbol stream.

Assuming the delay 640 is no... locations of the queues 126. This is because (1) the FIFOs 102 have been synchronized to operate in synchronism (a process described below), and (2) the push pointer counters 128 are clocked by the clock signal... of the symbol stream transmitted by the router 14A will be pulled from the clock synchronization FIFOs 102 of the CPUs 12A, 12B simultaneously, maintaining the required synchronization of received data when operating in duplex mode. In effect, the depths of the queues... order to achieve the operation just described with reference to Table 6, the reset and synchronization process shown in Fig 31A is used. The process not only initializes the clock synchronization FIFOs 102 of the CPUs 12A, 12B for duplex mode operation, but also operates to adjust the clock synchronization FIFOs 518 (Fig. 19A) of the CPU ports of each of

the routers 14A, 14B for duplex operation. The reset and synchronization process uses the SYNC command symbol to initiate a time period, delineated by the SYNC CLK signal 970 (Fig. 31B), to reset and initialize the respective clock synchronization FIFOs of the CPUs 12A and 12B and routers 14A, 14B. (The SYNC CLK signal... It is of a lower frequency than that used to receive symbols by the clock synchronization FIFOs, T(underscore)Clk. For example, where T(underscore)Clk is approximately 50 MHz, the... signal is approximately 3.125 MHz.)

Turning now to Fig. 31A, the reset and initialization process begins at step 950 by switching the clock signals used by the CPUs 12A, 12B and routers 14A, 14B as the transmit (T(underscore)Clk) and the unit's local clock (Local Clk) clock signals so that they are derived from the same... In addition, configuration registers in the CPUs 12A, 12B (configuration registers 74 in the interface units 24) and the routers 14A, 14B (contained in control logic unit 509 of routers 14A, 14B) are set to the FreqLock state.

The following discussion involves step 952, and makes reference to the interface unit 24 (Fig.5), router 14A (Fig. 19A) and Figs. 31A and 31B. With the clock... otherwise be sent followed by a self-addressed message packet. Any message packet in the process of being received and retransmitted when the SLEEP command symbols are received and recognized by... per the destination address). The SLEEP command symbol operates to "quiece" router 14A for the synchronization process. The self-addressed message packet sent by the CPU 12A, when received back by the... message packet sent after the SLEEP command symbol would necessarily have to be the last processed by the router 14A.

At step 954 the CPU 12A checks to see if it...the router will assert a RESET signal 972 that is applied to the two clock synchronization FIFOs 518 contained in the input logic 505(sub(4)), 505(sub(5)) of the... receive symbols directly from CPUs 12A, 12B. RESET, while asserted, will hold the two clock synchronization FIFOs 518 in a temporarily non-operating reset state with the push and pull pointer... As each of the CPUs 12 receive SYNC symbols are detected by the storage and processing units of the packet receivers 96 (Figs. 5 and 6) cause the RESET signal to be asserted by the packet receivers 96 (actually, storage and processing elements 110; Fig. 6) of each CPU 12. the RESET signal is applied to the... 4)), CPUs 12 and routers 14A, 14B de-assert the RESET signals, and the clock synchronization FIFOs of the CPUs 12A, 12, and routers 14A, 14B are released from their reset... the delay, the router 14A and CPUs 12 resume pulling data from their respective clock synchronization FIFOs and resume normal operation. The clock synchronization FIFOs of the router 14A begin pulling symbols from the queue (previously set by RESET... from the CPU 12A with the T(underscore)Clk will be pushed onto the clock synchronization FIFO at, for example, queue location 0 (or whatever other location pointed to by the... 0 (or whatever other location the push pointer was set to by RESET). The clock synchronization FIFOs of the router 14A are now synchronized to accommodate whatever delay 640 may be present in one communications path, relative to the... and the CPUs 12A, 12B.

Similarly, at the same virtual time, operation of the clock synchronization FIFOs 102 of both CPUs 12A, 12B is resumed, synchronizing them to the router 14A. Also, the CPUs 12A, 12B quit sending the SLEEP command... in favor of READY symbols, and resume message packet transmission, as appropriate.

That completes the synchronization process for the router 14A. However, the process must also be performed for the router 14B. Thus, the CPU 12A returns to step... however, assuming that the CPUs 12A, 12B are operating in duplex mode, the method and apparatus used to detect and handle a possible error, resulting in divergence of the CPUs from... via a message packet destined for a peripheral device of one or the other sub-processor systems 10A, 10B. Depending upon the destination of the outgoing message packet, step 1002 will... router 14 will issue an ERROR signal to the router control logic 509, causing the process to move to step 1004 where the router 14 detecting divergence will transmit a DVRG... time outs to occur. A router detecting divergence (without also detecting any simple link error) buys itself time to check the CRC of the received message packet by waiting for the... router 14, or received, all further message packets received from the CPUs and in the process of being routed

when divergence was detected, or the DVRG symbol received, will be passed... 1010) contained in one of the configuration registers 74 (Fig. 5) of the interface unit 24 of each CPU.

Returning for the moment to step 1006, the determination of which... "local" is meant to refer to the router 14A, 14B contained in the same sub-processor system 10A, 10B as the CPU. For example, referring to Fig. 1A, router 14A is... bit mentioned above: the bit contained in one of the configuration registers 74 of interface unit 24 (Fig. 5) of each CPU. When set to a first state, that particular CPU... the other CPU. In response, the state machines (not shown) within the control and status unit 509 (Fig. 19A) changes the "favorite" bits described above.

A few examples may facilitate understanding... DVRG symbol will echo that symbol to the routers 14A, 14B, start its internal divergence process timer, and begin determination of whether to continue or terminate. Having received a TLB symbol... to diverge with no errors reported. This can happen only if software (running on the processors 20) uses known divergent data to alter state. For example, suppose each CPU 12 has... number of the CPU 12A will differ from that of the CPU 12B. If the processors use the serial number to change the sequence of instructions executed (say, by branching if the serial number comes after some value) or to modify the value contained in a processor register, the complete "state" of the CPUs 12 will differ. In such cases, the "asymmetrical... of the primary CPU simply allows one CPU, and thereby the system 10, to continue processing without software intervention.

- An error at the output of the interface unit 24 of a CPU 12 will be detected by the router 14A, 14B, depending upon... router 14A, 14B that connects to a CPU 12 will be detected by the interface unit 24 of the affected CPU. The CPU will send a TLB symbol to the faulty... possible failure and, without external intervention, and transparently to the system user, remove the failing unit (CPU 12A or 12B, or router 14A or 14B) from the system to obviate or... reintegration." The discussion will refer to the CPUs 12A, 12B, routers 14A, 14B, and maintenance processor 18A, 18B shown forming parts of the processing system 10 illustrated in Fig. 1A. In addition, discussion will refer to the processors 20a, 20b, the interface units 24a, 24b, and the memory controllers 26a, 26b (Fig. 2) of the CPUs 12A, 12B as single units, since that is the way they function.

Reintegration is used to place two CPUs in... both of the paired CPUs at virtually the same time.

The major steps in the process of changing from simplex mode operation of the one on-line CPU to duplex mode... greater detail by the flow diagrams of Figs. 33A - 33D, generally are:

1. Setup and synchronize the two CPUs (one on-line, the other off-line) and their connected routers to... the memory of the on-line CPU to the off-line CPU, maintaining a tracking process that monitors changes in the memory of the on-line CPU that have not been... and may need to be copied over to, the off-line CPU;
3. Setup and synchronize the CPUs to run a delayed (slave) duplex mode from the same instruction stream (lock... will write the predetermined registers (not shown) of the control registers 74 in the interface units 24 of CPUs 12A and 12B, to a next state (after a soft operation) in... the off-line CPU 12B.

Next, a sequence is entered (steps 1060 - 1070) that will synchronize the clock synchronization FIFOs of the CPUs 12A, 12B and routers 14A, 14B in much the same fashion... the same steps described above in connection with the discussion of Figs. 31A, 31B to synchronize the clock synchronization FIFOs. The on-line CPU 12A will send the sequence of a SLEEP symbol, self-addressed message packet, and SYNC symbol which, with the SYNC CLK signal, operates to synchronize CPUs and routers. Once so synchronized, the on-line CPU 12A then, at step 1066, sends a Soft Reset (SRST) command... of all configuration registers and control registers (e.g., configuration registers 74 of the interface units 24) cache, and the like to memory 28 of the on-line CPU, copying... time to have the system 10 off-line for reintegration. For that reason, the reintegration process is performed in a manner that allows the on-line CPU to continue executing user... not match that of the off-line CPU. The reason for this is that normal processing by the processor 20 of the on-line CPU can change memory content after it has been copied... when a memory location is written in the on-line CPU 12A during the reintegration process it is marked as "dirty;"

second, all copying of memory to the off-line CPU... may, however, limit the ability to detect two-bit errors. But, since the memory copying process will last for a only relatively short period of time, this risk is believed acceptable... memory location in CPU 12A is made (either an incoming I/O write, or a processor write operation). The returning data (that was copied over to the off-line CPU) would... controller 26 (Fig. 2) of the on-line CPU to monitor memory locations in the process of being copied over to the off-line CPU 12B. The memory controller uses a... within the block had been written by another operation (e.g., a write by the processor 20, an I/O write, etc.), that prior write operation will flag the location in... still must be copied over to the off-line CPU 12B.

Returning to the reintegration process, and now to Fig. 33B, the memory tracking (AtomicWrite mechanism and using ECC to mark... entails writing a reintegration register (not shown; one of the configuration registers 74 of interface unit 24 - Fig. 5) to cause a reintegration (REINT) signal to be asserted. The REINT signal... is left alone. Throughout the incremental copy operations, the normal actions of the on-line processor will mark some memory locations dirty.

Several passes of incremental copying will need to be... the number of successful WriteConditional operations at the end of each pass through memory, the processors 20 can determine the effect of a given pass compared to the previous pass. When the benefits drop off, the processors 20 will give up on the precopy operations. At this point the reintegration process is ready to place the two CPUs 12A, 12B into lock-step operation.

Thus, the... in Fig. 33C, where at step 1100, the on-line CPU 12A momentarily halts foreground processing, i.e., execution of a user application. The remaining state (e.g., configuration registers, cache, etc.) of the on-line processors 20 and its caches is then read and written to a buffer (series of memory... to the off-line CPU 12B, together with a "reset vector" that will direct the processor units 20 of both CPUs 12A, 12B to a reset instruction.

Next, step 1106 will quiesce... to ensure that the FIFOs of the routers are clear, that the FIFOs of the processor interfaces 24 are clear, and no further incoming I/O message packets are forthcoming. At... symbol will be received and acted upon by both CPUs 12A, 12B, to cause the processor units 20 of each CPU to jump to the location in memory 28 containing the reset... a subroutine that will restore the stored state of both CPUs 12A, 12B to the processor units 20, caches 22, registers, etc. The CPUs 12A, 12B will then begin executing the same... enabling of the ECC bit to mark dirty locations must now be disabled, since the processors are doing the same thing to the same memory. During this stage of the reintegration... encountered by CPU 12A.

Meanwhile, the bus error in the CPU 12A will cause the processor unit 20 to be forced into an error-handling routine to determine (1) the cause of... error was caused by an attempt to read a memory location marked dirty. Accordingly, the processor unit 20 will initiate (via the BTE 88 -- Fig. 5) the AtomicWrite mechanism to copy the... the SRST symbols are now received by the CPUs 12A, 12B, they will cause both processor units 20 of the CPUs to be reset to start from the same location with the... will periodically update, e.g., a database or audit file that is indicative of the processing of the primary CPU up to that point in time of the update. Should the... in error-checking redundancy to the CPU 12B, in the same manner that the individual processor units 20a, 20b of the CPU 12A provide fail-fast, fault tolerance for the CPU - when... cost system is applicable, as illustrated in Fig. 34. As shown in Fig. 34, a processing system 10' includes the CPU 12A and routers 14A, 14B structured as described above. The... and the CPUs are also the same.

Thus, the CPU 12B' comprises only a single processor unit 20' and associated support components, including the cache 22', interface unit (IU) 24', memory controller 26', and memory 28'. Thus, while the CPU 12A is structured in the manner shown in Fig. 2, with cache processor unit, interface unit, and memory control redundancies, approximately one-half of those components are needed to implement CPU... stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... i.e., parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of...
...inputs connected to receive the data output

from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the...
...attached, without the assistance of any other CPU 12 in the system. Many prior parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O device becomes... rest of the system. Other prior systems provide access to I/O through pairs of processors so that should one of the processors fail, access ...if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... multiprocessing system access to every peripheral, as done here, operates to extend the "primary"/"backup" process taught in the above-identified U.S. Patent No. 4,228,496. There, a multiple CPU system may have a primary process may running on one CPU, while a backup process resides in the background on another of the CPUs. Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be... ..e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system... CPU can access any I/O without the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time... indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Specification: ...the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be...e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system... CPU can access any I/O without the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time... indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Claims: ...A2

1. In a computing system having at least a pair of processor units operating to execute substantially instructions in synchronous fashion, each of the pair of processor units communicatively coupled to first and second data communicating elements that operate to communicate data to and from the pair of processor units, a method for fault tolerant operation of the pair of processor units comprising the steps of:

the first and second data communicating elements receiving and comparing the data from the pair of processor units and at least one of the first and second data communicating elements operating to transmit to the pair of processor units an error signal indicative of miscompare of the data;

each of the processor units receiving the error signal and transmitting in response an echoed error signal to the first and second data communicating elements;

each of the pair of processor units determining from the error signals and the echoed error signals whether to continue or not continue operation; and

one of the pair of processor units continuing operation, and the other of the pair of processor units terminating operation.

2. The method of claim 1, wherein each of the first and second data communicating elements not detecting miscompare of the data transmitting to the pair of processor units another echoed error signal upon receipt of the echoed error signal.

3. The method of claim 2, including the steps of the other of the pair of processor units, after receipt of the error signal or the another echoed error signal, detecting a data... of the first and second data communicating elements is coupled to at least one peripheral unit for communicating data to and from the peripheral unit and the pair of processing units, and including the step of continuing transmission of any data communicated from the pair of processor units to the peripheral unit, and holding a last portion of data to the peripheral unit if a miscompare of data is detected.

5. The method of claim 1, including the step of the one of the pair of processor units transmitting to the first and second data communicating elements an ownership signal.

6. The method... second data communicating elements to disregard data transmission from the other of the pair of processor units.

7. The method of claim 5, wherein the ownership signal instructs the first and second data communicating elements to communicate only with the one of the pair of processor units.

8. The method of claim 5, wherein the first data communicating element favours communications from the one of the pair of processor units, and the second data communicating element favours communications with the other of the pair of processor units, and wherein messages from the other of the pair of processor units being concluded for

transmission to the peripheral unit with a data indication of bad data; and from the one of the pair or processing units being concluded with another data indication of good data.

Claims: ...B1

1. Procédé de fonctionnement insensible aux défaillances de deux unités de processeur (12A, 12B) dans un système informatique, lesdites deux unités de processeur fonctionnant pour exécuter des instructions sensiblement identiques en synchronisation, chacune des deux unités de processeur étant couplée pour communication à des premier et deuxième éléments de communication de données (14A, 14B) qui fonctionnent pour échanger des données entre les deux unités de processeur, le procédé comprenant les étapes consistant en :

la réception et la comparaison (1002) par les... ..deuxième éléments de communication de données (14A, 14B) des données provenant des deux unités de processeur (12A, 12B) et le fonctionnement d'au moins l'un des premier et deuxième éléments de communication de données pour transmettre (1004) aux deux unités de

processeur un signal d'erreur (DVRG) indicatif d'une erreur de comparaison des données ;

la réception par chacune des unités de processeur du signal d'erreur et la transmission par celles-ci en réponse d'un signal... ..14A, 14B) ne détectant pas d'erreur de comparaison pour transmettre aux deux unités de processeur un autre signal d'erreur renvoyé en écho lors de la réception du signal d'erreur renvoyé en écho provenant d'au moins l'une des deux unités de processeur ;

la détermination (1006) par chacune des deux unités de processeur, par l'analyse des diverses indications d'erreur qui leur sont fournies, s'il convient... ..deux unités et l'arrêt du fonctionnement (1012) par l'autre des deux unités de processeur.

2. Procédé selon la revendication 1, comprenant les étapes consistant en la détection par l'autre des deux unités de processeur (12A, 12B), après réception du signal d'erreur ou de l'autre signal d'erreur... ..17, 18) pour échanger des données entre l'unité périphérique et les deux unités de processeur, et comprenant l'étape consistant à poursuivre la transmission de toutes données communiquées par les deux unités de processeur à l'unité périphérique et la conservation d'une dernière partie de données dans l... ..1, comprenant l'étape consistant en la transmission par ladite une des deux unités de processeur aux premier et deuxième éléments de communication de données (14A, 14B) sur les trajets de... ..14B) d'ignorer les transmissions de données provenant de l'autre des deux unités de processeur.

6. Procédé selon la revendication 4, dans lequel le signal de propriété (IOY) donne l... ..éléments de communication de données de communiquer uniquement avec ladite une des deux unités de processeur.

7. Procédé selon la revendication 4, dans lequel, lors de la réception des données de... ..de communication de données ignorent les communications provenant de l'autre des deux unités de processeur, et dans lequel les données provenant de l'autre des deux unités de processeur se termineront pour transmission à l'une des unités périphériques par une indication de données de données erronées (TPB) et les données provenant de ladite une des deux unités de processeur se termineront par une autre indication de données de données correctes (TPG).

8. Système informatique comprenant :

au moins deux unités de processeur (12A, 12B) adaptées chacune pour exécuter la même instruction de flux de données sensiblement identiques... ..de communication de données (14A, 14B) couples pour communication à chacune de deux unités de processeur pour échanger des données entre celles-ci, chacun des premier et deuxième éléments de communication... ..i) des moyens pour comparer les données reçues de l'une des deux unités de processeur avec les données provenant de l'autre des deux unités de processeur et pour détecter une erreur de comparaison de celles-ci ;

(ii) des moyens pour transmettre, aux deux unites de processeur, un signal d'erreur (DVRG) indicatif d'une quelconque erreur de comparaison des donnees ;

des moyens, dans chacune desdites deux unites de processeur, pour recevoir le signal d'erreur et pour transmettre en reponse un signal d'erreur... ...comparaison des donnees n'a pas ete detectee, fonctionner pour transmettre aux deux unites de processeur un autre signal d'erreur renvoye en echo lors de la reception du signal d'erreur renvoye en echo en provenance d'au moins l'une des deux unites de processeur ;

des moyens, dans chacune des deux unites de processeur, pour determiner, par l'analyse des diverses indications d'erreur qui leur sont fournies, s...

6/K/27 (Item 27 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Patent Assignee:

- Compaq Computer Corporation... ;
;;

Country	Number	Kind	Date	
Type	Pub. Date		Kind	Text
...Compaq Computer Corporation (687797) 20555 SH 249 Houston, Texas 77070-2698 US		19		
Examination...		19		
Available Text		Language	Update	Word Count
Total Word Count (Document A)				
Total Word Count (Document B)				
Total Word Count (All Documents)				

Specification: ...and 08/485,053, filed concurrently herewith.

The present invention is directed generally to data processing systems, and more particularly to a multiple processing system and a reliable system area network that provides connectivity for interprocessor and input/output... ...and communications systems to general purpose high availability commercial systems. The evolution of fault tolerant computers has been well documented (see D. P. Siewiorek, R. S. Swarz, "The Theory and Practice... ...and the Jet Propulsion laboratory began to apply fault tolerance to the development of guidance computers for aerospace applications. The 1960's also saw the development of the first AT&T electronic switching systems.

The first commercial fault tolerant machines were introduced by Tandem Computers in the 1970's for use in on-line transaction processing applications (J. Bartlett, "A NonStop Kernal," in proc. Eighth Symposium on Operating System Principles, pp... ...systems were introduced in the 1980's (O. Serlin, "Fault- Tolerant Systems in Commercial Applications," Computer, pp. 19-30, August 1984). Current commercial fault tolerant systems include distributed memory multi-processors, shared-memory transaction based systems, "pair-and-spare" hardware fault tolerant

systems (see R. Freiburghouse, "Making Processing Fail-safe," Mini-micro Systems, pp. 255-264, May 1982; U.S. Patent No. 4,228,496), and triple-modular-redundant systems such as the "Integrity" computing system manufactured by Tandem Computers Incorporated of Cupertino, California, assignee of this application and the invention disclosed herein.

Most applications of commercial fault tolerant computers fall into the category of on-line transaction processing. Financial institutions require high availability for electronic funds transfer, control of automatic teller machines, and... telecommunications systems.

Vendors of fault tolerant machines attempt to achieve both increased system availability, continuous processing, and correctness of data even in the presence of faults. Depending upon the particular system architecture, application software ("processes") running on the system either continue to run despite failures, or the processes are automatically restarted from a recent checkpoint when a fault is encountered. Some fault tolerant systems are provided with sufficient component redundancy to be able reconfigure around failed components, but processes running in the failed modules are lost. Vendors of commercial fault tolerant systems have extended fault tolerance beyond the processors and disks. To make large improvements in reliability, all sources of failure must be addressed... power supplies, fans and inter-module connections.

The "NonStop," and "Integrity" architectures manufactured by Tandem Computers Incorporated, (both respectively illustrated broadly in U.S. Patent No. 4,228,496 and U.S. Patent No. 4,278,496, both assigned to the assignee of this application; NonStop and Integrity are registered trademarks of Tandem Computers Incorporated) represent two current approaches to commercial fault tolerant computing. The NonStop system, as generally... above-identified U.S. Patent No. 4,278,496, employs an architecture that uses multiple processor systems designed to continue operation despite the failure of any single hardware component. In normal operation, each processor system uses its major components independently and concurrently, rather than as "hot backups". The NonStop system architecture may consist of up to 16 processor systems interconnected by a bus for interprocessor communication. Each processor system has its own memory which contains a copy of a message-based operating system. Each processor system controls one or more input/output (I/O) busses. Dual-porting of I/O controllers and devices provides multiple paths to each device. External storage (to the processor system), such as disk storage, may be mirrored to maintain redundant permanent data storage.

This... hardware, while fault recovery is the responsibility of the software.

Also, in the Nonstop multi-processor architecture, application software ("process") may run on the system under the operating system as "process-pairs," including a primary process and a backup process. The primary process runs on one of the multiple processors while the backup process runs on a different processor. The backup process is usually dormant, but periodically updates its state in response to checkpoint messages from the primary process. The content of a checkpoint message can take the form of complete state update, or currently most application code runs under transaction processing software which provides recovery through a combination of checkpoints and transaction two-phase commit protocols.

Interprocessor message traffic in the Tandem Nonstop architecture includes each processor periodically broadcasting an "I'm Alive" message for receipt by all the processors of the system, including itself, informing the other processors that the broadcasting processor is still functioning. When a processor fails, that failure will be announced and identified by the absence of the failed processor's periodic "I'm Alive" message. In response, the operating system will direct the appropriate backup processes to begin primary execution from the last checkpoint. New backup processes may be started in another processor, or the process may be run with no backup until the hardware has been repaired. U.S. Patent... example of this technique.

Each I/O controller is managed by one of the two processors to which it is attached. Management of the controller is periodically switched between the processors. If the managing processor fails, ownership of the controller is

automatically switched to the other processor. If the controller fails, access to the data is maintained through another controller.

In addition to providing hardware fault tolerance, the processor pairs of the above-described architecture provide some measure of software fault tolerance. When a processor fails due to a software error, the backup processor frequently is able to successfully continue processing without encountering the same error. The software environment in the backup processor typically has different queue lengths, table sizes, and process mixes. Since most of the software bugs escaping the software quality assurance tests involve infrequent data dependent boundary conditions, the backup processes often succeed.

In contrast to the above-described architecture, the Integrity system illustrates another approach... fault recovery is the logical choice since few modifications to the software are required. The processors and local memories are configured using triple-modular-redundancy (TMR). All processors run the same code stream, but clocking of each module is independent to provide tolerance... three streams is asynchronous, and may drift several clock periods apart. The streams are re-synchronized periodically and during access of global memory. Voters on the TMR Controller boards detect and mask failures in a processor module. Memory is partitioned between the local memory on the triplicated processor boards and the global memory on the duplicated TMRC boards. The duplicated portions of the... techniques to detect failures. Each global memory is dual ported and is interfaced to the processors as well to the I/O Processors (IOPs). Standard VME peripheral controllers are interfaced to a pair of busses through a Bus... the BIMs to switch control of all controllers to the remaining IOP. Mirrored disk storage units may be attached to two different VME controllers.

In the Integrity system all hardware failures... reintegrated on-line.

The preceding examples illustrate present approaches to incorporating fault tolerance into data processing systems. Approaches involving software recovery require less redundant hardware, and offer the potential for some... have been developed on other systems.

Thus, the systems described above provide fault tolerant data processing either by hardware (e.g., fail-functional, employing redundancy) or by software techniques (fail-fast... hardware). However, none of the systems described are believed capable of providing fault tolerant data processing, using both hardware (fail-functional) and software (fail-fast) approaches, by a single data processing system.

Computing systems, such as those described above, are often used for electronic commerce: electronic data interchange (EDI) and global messaging. Today's demands upon such electronic commerce, however, is demanding more and more throughput capacity as the number of users increases and... networks such as local area networks (LANS), and the like.

A key requirement for a server architecture is the ability to move massive quantities of data. The server should have high bandwidth that is scalable, so that added throughput capacity can be added... response time, latency affects service levels and employee productivity.

The present invention provides a multiple-processor system that combines both of the two above-described approaches to fault tolerant architecture, hardware redundancy and software recovery techniques, in a single system.

Broadly, the present invention includes a processing system composed of multiple sub-processing systems. Each sub-processing system has, as the main processing element, a central processing unit (CPU) that in turn comprises a pair of processors operating in lock-step, synchronized fashion... execute each instruction of an instruction stream at the same time. Each of the sub-processing systems further include an input/output (I/O) system area network system that provides redundant communication paths between various components of the larger processing

system, including a CPU and assorted peripheral devices (e.g., mass storage

units, printers, and the like) of a sub-processing system, as well as between the sub-processors that may make up the larger overall processing system. Communication between any component of the processing system (e.g., a CPU and another CPU, or a CPU and any peripheral device, regardless of which sub-processing system it may belong to) is implemented by forming and transmitting packetized messages that are... responsible for choosing the proper or available communication paths from a transmitting component of the processing system to a destination component based upon information contained in the message packet. Thus, the... peripherals, but permits it to also be used for interprocessor communications.

As indicated above, the processing system of the present invention is structured to provide fault-tolerant operation through both "fail... at a variety of points in the various data paths between the (lock-step operated) processor elements of the CPU and its associated memory. In particular, the processing system of the present invention conducts error-checking at an interface, and in a manner... little impact on performance. Prior art systems typically implement error-checking by running pairs of processors, and checking (comparing) the data and instruction flow between the processors and a cache memory. This technique of error-checking tended to add delay to the... error-checking precluded use of off-the-shelf parts that may be available (i.e., processor/cache memory combinations on a single semiconductor chip or module). The present invention performs error-checking of the processors at points that operate at slower rates, such as the main memory and I/O interfaces which operate at slower speeds than the processor-cache interface. In addition, the error-checking is performed at locations that allow detection of errors that may occur in the processors, their cache memory, and the I/O and memory interfaces. This allows simpler designs for... other data integrity checks.

Error-checking of the communication flow between the components of the processing system is achieved by adding a cyclic-redundancy-check (CRC) to the message packets that... "Good" (TPG) or "This Packet Bad" (TPB) - is appended to every packet. A maintenance diagnostic processor can use this information to isolate a link or router element that introduces an error... of topologies, so that alternate paths can be provided between any two elements of a processing system (e.g., between a CPU and an I/O device), for communication in the... so (e.g., by creating a "deadlock" condition, discussed further below).

The CPUs of a processing system are capable of operating in one of two basic modes: a "simplex mode" in... independently of the other, or a "duplex" mode in which pairs of CPUs operate in synchronized, lock-step fashion. Simplex mode operation provides the capability of recovering from faults that are... U.S. Pat. No. 4,228,496 which teaches a multiprocessing system in which each processor has the capability of checking on the operability of its sibling processors, and of taking over the processing of a processor found or believed to have failed). When operating in duplex mode, the paired CPUs both... fault tolerant platform for less robust operating systems (e.g., the UNIX operating system). The processing system of the present invention, with the paired, lock-step CPUs, is structured so that... masked (i.e., operating despite the existence of a fault), primarily through hardware.

When the processing system is operating in duplex mode, each CPU pair uses the I/O system to access any peripheral of the processing system, regardless of which (of the two, or more) sub-processor system the peripheral may be ostensibly a member of. Also, in duplex mode, message packets... message for the CPU pair (from either a peripheral device such as a mass storage unit or from a processing unit), will replicate the message and deliver it to both CPUs of the pair using synchronization methods that ensure that the CPUs remain synchronized. In effect, the duplex CPU pair, as viewed from the I/O system and other... as a single CPU. Thus, the I/O system, which includes elements from all sub-processing systems, is made to be seen by the duplex CPU pair as one homogeneous system... a multiprocessor system in which the CPU of any one is actually a pair of synchronized, lock-step CPUs.

Yet another important aspect of the present invention is that interrupts issuing... interrupts via the message packet system ensures that they will arrive at duplexed CPUs in synchronized fashion, in the same manner as I/O message packets. Interrupt message packets will contain... the system. In addition, using the same messaging system to communicate data between I/O units and the CPUs and to communicate interrupts to the CPUs preserves the ordering of I... the implementation of a technique of validating access to the memory of any CPU. The processing

system, as structured according to the present invention, permits the memory of any CPU to a CPU and any other component of the processor system. Thereby, the individual processor units of the CPU are removed from the more mundane tasks of getting information from memory and out onto the TNet network, or accepting information from the network. The processor unit of the CPU merely sets up data structures in memory containing the data to be... ..is required, where in memory the response is to be placed when received. When the processor unit completes the task of creating the data structure, the block transfer engine is notified to... ..response is received, it is routed to the expected memory location identified, and notifies the processor unit that the response was received.

Further aspects and features of the present invention will become... ..invention, which should be taken in conjunction with the accompanying drawings.

Fig. 1A illustrates a processing system constructed in accordance with the teachings of the present invention, and Figs. 1B and 1C illustrate two alternate configurations of the processing system of Fig. 1A, employing clusters or arrangements of the processing system of Fig. 1A;

Fig. 2 illustrates, in simplified block diagram form, the central processing unit (CPU) that forms a part of each sub-processor system of Figs. 1A - 1C;

Figs. 3A - 3D and 4A - 4C illustrate the construction of... ..the area network I/O system shown in Fig. 2;

Fig. 5 illustrates the interface unit that forms a part of the CPUs of Fig. 2 to interface the processor and memory with the I/O area network system;

Fig. 6 is a block diagram, illustrating a portion of packet receiver of the interface unit of Fig. 5;

Fig. 7A diagrammatically illustrates the clock synchronization FIFO (CS FIFO) used by the packet receiver section packet receiver shown in Fig. 6;

Fig. 7B is an block diagram of a construction of the clock synchronization FIFO structure shown in Fig. 7A;

Fig. 8 illustrates the cross-connections for error-checking outbound transmissions from the two interface units of a CPU;

Fig. 9 illustrates an encoded (8B to 9B) data/command symbol;

Fig. 10 illustrates the method and structure used by the interface unit of Fig. 5 to cross-check for errors data being transferred to the memory controllers... ..of a CPU of Fig. 2 to other (external to the CPU) components of the processing system;

Fig. 12 is a block diagram that diagrammatically illustrates the formation of an address... ..14A illustrates the logic for posting interrupt requests to queues in memory and to the processor units of the CPU of Fig. 2;

Fig. 14B illustrates the process used to form a memory address for a queue entry;

Fig. 15 is a block... ..data output constructs formed in the memory of the CPU of Fig. 2 by a processor unit, and containing data to be sent via the area I/O networks shown in Figs. 1A - 1C, and also illustrating the block transfer engine (BTE) unit of the interface unit of Fig. 5 that operates to access the data output constructs for transmission to the... ..pair of memory controllers between memory of a CPU of Fig. 2 and its interface unit for accessing from memory 72 bits of data, including two simultaneously-accessed 32-bit words... ..other for error-checking;

Fig. 19A is a simplified block diagram illustration of the router unit used in the area input/output networks of the processing systems shown in Figs. 1A - 1C;

Fig. 19B illustrates comparison on two port inputs of the router unit of Fig. 19A;

Fig. 20A is a block diagram the construction of one of the six input ports of the router unit shown in Fig. 19A;

Fig. 20B is a block diagram of the synchronization logic used to validate command/data symbols received at an input port of the router unit of Fig. 19A;

Fig. 21A is a block diagram illustration of the target port selection... ..is a block diagram illustration of one of the six output ports of the router unit shown in Fig. 19A;

Fig. 23 is an illustration of the method used to transmit identical information to a duplexed pair CPUs of Fig. 2 in synchronized fashion when the processing system is operating in lock-step (duplex) mode, using a pair the FIFOs of Fig... ..is a simplified block diagram illustrating the clock generation system of each of the sub-processing systems of Figs. 1A - 1C for developing the plurality of clock signals used to operate the various elements of that sub-processing system;

Fig. 25 illustrates the topology used to interconnect the clock generation systems of paired sub-processing systems for synchronizing the various clock signals of the pair of sub-processing systems to one another;

Fig. 26A and 26B illustrates a FIFO constant rate clock control logic used to control the clock synchronization FIFO of Figs. 8 or 20 in the situation when the two clocks used to... ..structure of the on-line access port (OLAP) used to provide access to the maintenance

processor (MP) to the various elements of the system of Fig. 1A (or those of Figs... ..the soft-flag logic used to handle asymmetric variables between the CPUs of paired sub-processing systems operating in duplex mode;

Fig. 31A shows a flow diagram, and Fig. 31B illustrates a portion of SYNC CLK, both of which are used to reset and synchronize the clock synchronization FIFOs of the CPUs and routers of the processing system of Fig. 1A that receive information from each other;

Fig. 32 is a flow... ..33A - 33D generally illustrate the procedure used to bring an one of the CPUs of processing system shown in Fig. 1A into lock-step, duplex mode operation with the other of the CPUs without measurably halting operation of the processing system; and

Fig. 34 illustrates a reduced cost architecture incorporating teachings of the invention; and... ..to the figures and, for the moment, principally Fig. 1A, there is illustrated a data processing system, designated with the reference 10, constructed according to the various teachings of the present invention. As Fig. 1A shows, the data processing system 10 comprises two sub-processor systems 10A and 10B each of which are substantially the same in structure and function... ..should be appreciated that, unless noted otherwise, a description of any one of the sub-processor systems 10 will apply equally to any other sub-processor system 10.

Continuing with Fig. 1A therefore, each of the sub-processor systems 10A, 10B is illustrated as including a central processing unit (CPU) 12, a router 14, and a plurality of input/output (I/O) packet interfaces... ..one of the I/O packet interfaces 16 will also have coupled thereto a maintenance processor (MP) 18.

The MP 18 of each sub-processor system 10A, 10B connects to each of the elements of that sub-processor system via an IEEE 1149.1 test bus 17 (shown in phantom in Fig. 1A... ..accompanying clock signal. As Fig. 1A further illustrates, TNet Links L also interconnect the sub-processor systems 10A and 10B to one another, providing each sub-processor system 10 with access to the I/O devices of the other as well as inter-CPU communication. As will be seen, any CPU 12 of the processing system 10 can be given access to the memory of any other CPU 12, although... ..the memory of a CPU 12 by a wayward peripheral device 17.

Preferably, the sub-processor systems 10A/10B are paired as illustrated in Fig. 1A (and Figs 1B and 1C, discussed below), and each sub-processor system 10A/10B pair (i.e., comprising a CPU 12, at least one router 14... ..12A) connects, by a TNet Link L to a router (14A) of the corresponding sub-processor system (e.g., 10A). Conversely, the Y port connects the CPU (12A) to the router (14B) of the companion sub- processor system (10B). This latter

connection not only provides a communication path for access by a CPU (12A) to the I/O devices of the other sub-processor system (10B), but also to the CPU (12B) of that system for inter-CPU communication.

Information is communicated between any element of the processing system 10 and any other element (e.g., CPU 12A of sub-processor system 10A) of the system and any other element of the system (e.g., an I/O device associated with an I/O packet interface 16B of sub-processor system 10B) via message "packets." Each message packet is made up of a number of... this reason, a unique method of receiving the symbols at the receiver, using a clock synchronization first-in-first-out (CS FIFO) storage structure (described more fully below), has been developed... operation means just that: the frequencies of the clock signals of the transmitter and receiver units are locked, although not necessarily in phase. Frequency locked clock signals are used to transmit symbols between the routers 14A, 14B and the CPUs 12 of paired sub-processor systems (e.g., sub-processor systems 10A, 10B, Fig. 1A). Since the clocks of the transmitting and receiving element are not phase related, a clock synchronization FIFO is again used -- albeit operating in a slightly different mode from that used for... difference, as will be seen, is due to the fact that pairs of the sub-processor systems 10 can be operated in a synchronized, lock-step mode, called duplex mode, in which each CPU 12 operates to execute the... 1A illustrates another feature of the invention: a cross-link connection between the two sub-processor systems 10A, 10B through the use of additional routers 14 (identified in Fig. 1A as... RY(sub(1)), and RY(sub(2)) form a cross-link connection between the sub-processors 10A, 10B (or, as shown, "sides" X and Y, respectively) to couple them to I... the routers RX(sub(2)) and RY(sub(2)) provide the I/O packet interface units 16x and 16y with a dual ported interface. Of course, it will now be evident... lend themselves to being used in a manner that can extend the configuration of the processing system 10 to include additional sub-processor systems such as illustrated in Figs. 1B and 1C. In Fig. 1B, for example, one... of each of the routers 14A and 14B is used to connect the corresponding sub-processor systems 10A and 10B to additional sub-processor systems 10A' and 10B' forming thereby a larger processing system comprising clusters of the basic processing system 10 of Fig. 1.

Similarly, in Fig. 1C the above concept is extended to form an eight sub-processor system cluster, comprising sub-processor systems pairs 10A/10B, 10A'/10B', 10A''/10B'', and 10A'''/10B'''. In turn, each of the sub-processor systems (e.g., sub-processor system 10A) will have essentially the same basic minimum configuration of a CPU 12, a... by a I/O packet interface 16, except that, as Fig. 1C shows, the sub-processor systems 10A and 10B include additional routers 14C and 14D, respectively, in order to extend the cluster beyond sub-processor systems 10A'/10B' to the sub-processor systems 10A''/10B'' and 10A'''/10B'''. As Fig. 1C further illustrates, unused ports 4 and... the routers 14 when configuring the topology of the system 10, any CPU 12 of processing system 10 of Fig. 1C can access any other "end unit" (e.g., a CPU or I/O device) of any of the other sub-processor systems. Two paths are available from any CPU 12 to the last router 14 connecting to the I/O packet interface 16. For example, the CPU 12B of the sub-processor system 10B' can access the I/O 16''' of sub-processor system 10A''' via router 14B (of sub-processor system 10B'), router 14D, and router 14B (of sub-system 10B'') and, via link LA... 10A'''), OR via router 14A (of sub-system 10A'), router 14C, and router 14A (sub-processor system 10A'''). Similarly, CPU 12A of sub-processor system 10A'' may access (via two paths) memory contained in the CPU 12B of sub-processor 10B to read or write data. (Memory accesses by one CPU 12 of another component of the processing system requires, as will be seen, the components seeking access to have authorization to do... prevents corruption of memory data of a CPU by erroneous access.)

The topology of the processing system shown in Fig. 1B is achieved by using port 1 of the routers 14A, 14B, and auxiliary TNet links LA, to connect to the routers 14A', 14B' of sub-processor systems 10A', 10B'. The topology thereby obtained establishes redundant communication paths between any CPU 12 (12A, 12B, 12A', 12B') and any I/O packet interface 16 of the processing system 10 shown in Fig. 1B. For example, the CPU 12A' of the sub-processor system 10A' may access the I/O 16A of sub-processor system 10A by a first path formed by the router 14A' (in port 4, out... shown in Fig. 1B. By interconnecting one port of each router 14 of each sub-processor pair, and using additional auxiliary TNet links LA (illustrated in Fig. 1C with the dotted line connections) between the ports 1 of the routers 14 (14A'' and 14B'') of sub-processor systems 10A'', 10B'' and 10A''', 10B''', two separate,

independent data paths can be found between any CPU 12 and any I/O packet interface 16. In this fashion, any end unit (i.e., a CPU 12 or an I/O packet interface 16) will have at least two paths to any other end unit.

Providing alternate paths of access between any two end units (e.g., between a CPU 12 and any other CPU 12, or between any CPU... any two of the remaining fault domains. Here, a fault domain could be a sub-processor system (e.g., 10A). Thus, if the sub-processor system 10A were brought down because of a failure the electrical power being supplied, without... TNet link LA between the routers 14A' and 14B', the CPU 12B of the sub-processor system 10B would have lost access to the I/O packet interface 16' (via router... with the loss of the router 14A (and router 14C) by loss of the sub-processor system 10A, communications between the CPU 12B is still possible via the route of router... equally to CPU 12B. As Fig. 2 shows, the CPU 12A includes a pair of processor units 20a, 20b that are configured for synchronized, lock-step operation in that both processor units 20a, 20b receive and execute identical instructions, and issue identical data and command outputs, at substantially the same moments in time. Each of the processor units 20a and 20b is connected, by a bus 21 (21a, 21b) to a corresponding cache memory 22. The particular type of processor units used could contain sufficient internal cache memory so that the cache memory 22 would not... 22 could be used to supplement any cache memory that may be internal to the processor units 20. In any event, if the cache memory 22 is used, the bus 21 is... 22 address bits, 3 bits of parity covering the address, and 7 control bits.

The processors 20a, 20b are also respectively coupled, via a separate 64-bit address/data bus 23 to X and Y interface units 24a, 24b. If desired, the address/data communicated on each bus 23a, 23b could also be protected by parity, although this will increase the width of the bus. (Preferably, the processors 20 are constructed to include RISC R4000 type microprocessors, such as are available from the MIPS Division of Silicon Graphics, Inc. of Santa Clara, California.)

The X and Y interface units 24a, 24b operate to communicate data and command signals between the processor units 20a, 20b and a memory system of the CPU 12A, comprising a memory controller (MC... MC halves 26a and 26b) and a dynamic random access memory array 28. The interface units 24 interconnect to each other and to the Mcs 26a, 26b by a 72-bit... accompanied by 8 bits of ECC) are written to the memory 28 by the interface units 24, one interface unit 24 will drive only one word (e.g., the 32 most significant portion) of the doubleword being written while the other interface unit 24 writes the other word of the double word (e.g., the least significant 32-bit portion of the doubleword). In addition, on each write operation the interface units 24a, 24b perform a cross-check operation on the data not written by that interface unit 24 with the data written by the other to check for errors; on read operations... accessed corresponds to the address of the location from which the doubleword was stored.

Interface units 24a, 24b of the CPU 12A form the circuitry to respectively service the X and Y (I/O) ports of the CPU 12A. Thus, the X interface unit 24a connects by the bi-directional TNet Link Lx to a port of the router 14A of the processor system 10A (Fig. 1A) while the Y interface unit 24b similarly connects to the router 14B of the processor system 10B by TNet Link Ly. The X interface unit 24a handles all I/O traffic between the router 14A and the CPU 12A of the sub-processor system 10A. Likewise, the Y interface unit 24b is responsible for all I/O traffic between the CPU 12A and the router 14B of companion sub-processor system 10B.

The TNet Link Lx connecting the X interface unit 24a to the router 14A (Fig. 1) comprises, as above indicated, two 10-bit buses... sub(x)) carries data incoming from the router 14A. In similar fashion, the Y interface unit 24b is connected to the router 14B (of the sub-processor system 10B) by two 10-bit busses: 30(sub(y)) (for outgoing transmissions) and 32... y)) (for incoming transmissions), together forming the TNet Link Ly.

The X and Y interface units 24a, 24b are synchronously operated in lock-step, performing substantially the same operations at substantially the same times. Thus, although only the X interface unit 24a actually transmits data onto the bus 30(sub(x)), the same output data is being produced by the Y interface unit 24b, and used for error-checking. The Y interface unit 24b output data is coupled to the X interface unit 24a by a cross-link 34(sub(y)) where it is

received by the X interface unit 24a and compared against the same output data produced by the X interface unit. In this way the outgoing data made available at the X port of the CPU... the port of the CPU 12A is checked. The output data from the Y interface unit 24b is coupled to the Y port by a 10-bit bus 30(sub(y)), and also to the X interface unit 24a by the 9-bit cross-link 34(sub(y)) where is checked with that produced by the X interface unit.

As mentioned, the two interface units 24a, 24b operate in synchronous, lock-step with one another, each performing substantially the same... X and/or Y ports of the CPU 12A must be received by both interface units 24a, 24b to maintain the two interface units in this lock-step mode. Thus, data received by one interface unit 24a, 24b is passed to the other, as indicated by the dotted lines and 9... sub(x)) (communicating incoming data being received at the X port by the X interface unit 24a to the Y interface unit 24b) and 36(sub(y)) (communicating data received at the Y port by the Y interface unit 24b to the X interface unit 24a).

Certain more robust operating systems are structured with a fault-tolerant capability in the... example, U.S. Patent No. 4,817,091 teaches a multiprocessor system in which each processor periodically messages each of the processors of the system (including itself), under software control, to thereby provide an indication of continuing operation. Each of the processors, in addition to performing its normal tasks, operates as a backup processor to another of the processors. In the event one of the backup processors fails to receive the messaged indication from a sibling processor, it will take over the operation of that sibling (now thought to be inoperative), in... platform for both types of software. Thus, when a robust operating system is available, the processing system 10 can be configured to operate in a "simplex" mode in which each of... left, in most instances, to software.

Alternatively, for less robust operating systems and software, the processing system 10 provides a hardware-based fault-tolerance by being configured to operate in a... g., CPUs 12A, 12B) are coupled together as shown in Fig. 1A, to operate in synchronized, lock-step fashion, executing the same instructions at the substantially the same moment in time... data and command symbols. In order to simplify the design of the CPU 12, the processors 20 are precluded from communicating directly with any outside entity (e.g., another CPU 12... O device via the I/O packet interface 16). Rather, as will be seen, the processor will construct a data structure in memory and turn over control to the interface units 24. Each interface unit 24 includes a block transfer engine (BTE; Fig. 5) configured to provide a form of... to the destination according to information contained in the message packet.

The design of the processing system 10 permits a memory 28 of a CPU to be read or written by... via the routers 14. Accordingly, before continuing with the description of the construction of the processing system 10, it would be of advantage to understand first the configuration of the data... information.

As indicated, the HADC message packet operates to communicate write data between the end units (e.g., CPU 12) of the processing system 10. Other message packets, however, may be differently constructed because of their function and... CRC. The HC message packet is used to acknowledge a request to write data.

Interface Unit:

The X and Y interface units 24 (i.e., 24a and 24b - Fig. 2) operate to perform three major functions within the CPU 12: to interface the processors 20 to the memory 28; to provide an I/O service that operates transparently to, but under the control of, the processors; and to validate requests for access to the memory 28 from outside sources.

Regarding first the interface function, the X and Y interface units 24a, 24b operate to respectively communicate processors 20a, 20b to the memory controllers (Mcs 26a, 26b) and memory 28 for writing and... fast checking of the data read/written. For example, write operations have the two interface units 24a, 24b cooperating to cross-check the data to be written to ensure its integrity (and at the same time, the interface units 24 will operate) to develop an error correcting code (ECC) that covers, as will be appropriate address.

With respect to I/O access, the processors 20 are not provided with the ability to communicate directly with the input/output systems... must write data structures to the memory 28 and then pass control to the interface units 24

which perform a direct memory access (DMA) operation to retrieve those data structures, and... ..indicated in the data structure itself.)

The third function of the X and Y interface units 24, access validation to the memory 28, uses an address validation and translation (AVT) table maintained by the interface units. The AVT table contains an address for each system component (e.g., an I/O... ..the incoming message packets are virtual addresses. These virtual addresses are translated by the interface unit to physical addresses recognizable by the memory control units 26 for accessing the memory 28.

Referring to Fig. 5, illustrated is a simplified block diagram of the X interface unit 24a of the CPU 12A. The companion Y interface unit 24b (as well as the interface units 24 of the CPU 12B, or any other CPU 12) is of substantially identical construction. Accordingly, it will be understood that a description of the interface unit 24a will apply equally to the other interface units 24 of the processing system 10.

As Fig. 5 illustrates, the X interface unit 24a includes a processor interface 60, a memory interface 70, interrupt logic 86, a block transfer engine (BTE) 88, access validation and translation logic 90, a packet transmitter 94, and a packet receiver 96.

Processor Interface:

The processor interface 60 handles the information flow (data and commands) between the processor 20a and the X interface unit 24a. A processor bus 23, including a 64 bit address and data bus (SysAD) 23a and a 9 bit command bus 23b, couples the processor 20a and the processor interface 60 to one another. While the SysAD bus 23a carries memory address and data... ..and qualifying commands carried at substantially the same time on the SysAD bus 23a. The processor interface 60 operates to interpret commands issued by the processor unit 20a in order to pass reads/writes to memory or control registers of the processor interface. In addition, the processor interface 60 contains temporary storage (not shown) for buffering addresses and data for access to... ..26). Data and command information read from memory is similarly buffered en route to the processor unit 20a, and made available when the processor unit is ready to accept it. Further, the processor interface 60 will operate to generate the necessary interrupt signalling for the X interface unit 24a.

The processor interface 60 is connected to a memory interface 70 and to configuration registers 74 by a bi-directional 64 bit processor address/data bus 76. The configuration registers 74 are a symbolic representation of the various control registers contained in other components of the X interface unit 24a, and will be discussed when those particular components are discussed. However, although not specifically... ..throughout other of the logic that is used to implement the X interface 24a, the

processor address/data bus 76 is likewise coupled to read or write to those registers.

Configuration registers 74 are read/write accessible to the processor 20a; they allow the X interface unit to be "personalized." For example, one register identifies the node address of the CPU 12A... ..with the CPU 12A; another, readable only, contains a fixed identification number of the interface unit 24, and still other registers define areas of memory that can be used by, for... ..logic 90, etc.) employing them are discussed.

The memory interface 70 couples the X interface unit 24a to the memory controllers 26 (and to the Y interface unit 24b; see fig. 2) by a bus 25 that includes two 36 bi-directional bit... ..25a, 25b. The memory interface operates to arbitrate between requests for memory access from the processor unit 20, the BTE 88, and the AVT logic 90. In addition to memory accesses from the processor unit 20a, the memory 28 may also be accessed by components of the processing system 10 to, for example, store data requested to be read by the processor unit 20a from an I/O unit 17, or memory 28 may also be accessed for I/O data structures previously set up in memory by the processor unit. Since these accesses are all asynchronous, they must be arbitrated, and the memory interface 70... ..command

information accessed from the memory 28 is coupled from the memory interface to the processor interface 60 by a memory read bus 82, as well as to an interrupt logic... ..doubleword quantities. However, while the memory interfaces 70 of both the X and Y interface units 24a ...by the memory interface 70 are coupled to the memory interface by the companion interface unit 24 where they are compared with the same 32 bits for error.

Digressing for the... ..containing interrupt information are received, that information is conveyed to the interrupt logic 86 for processing and posting for action by the processor 20, along with any interrupts generated internal to the CPU 12A. Internally generated interrupts will... ..register 71 (internal to the interrupt logic 86), indicating the cause of the interrupt. The processor 20 can then read and act upon the interrupt. The interrupt logic is discussed more fully below.

The BTE 88 of the X interface unit 24a operates to perform direct memory accesses, and provides the mechanism that allows the processors 20 to access external resources. The BTE 88 can be set-up by the processors 20 to generate I/O requests, transparent to the processors 20 and notify the processors when the requests are complete. The BTE logic 88 is discussed further below.

Requests for... ..8 byte wide format necessary for storing in the memory 28.

Outgoing message packets containing processor originated transaction requests (e.g., a read request asking for a block data from an I/O unit) are monitored by the request transaction logic (RTL) 100. The RTL 100 provides a time... ..will generate an interrupt (handled and reported by the interrupt logic 86) to inform the processor 20 that the request was not honored. In addition, the RTL 100 will validate responses... ..28 (by the DMA operation of the BTE 86) at a location known to the processor 20 so that it can locate the response.

Each of the CPUs 12 are checked... ..discussed. One such check is an on-going monitor of the operation of the interface units 24a, 24b of each CPU. Since the interface units 24a, 24b operate in lock-step synchronism checking can be performed by monitoring the operating states of the paired interface units 24a, 24b by a continuous comparison of certain of their internal states. This approach is implemented by using one stage of a state machine (not shown) contained in the unit 24a of CPU 12A, and comparing each state assumed by that stage with its identical state machine stage in the interface unit 24b. All units of the interface units 24 use state machines to control their operations. Preferably, therefore, a state machine of the memory interface 70 that controls the data transfers between the interface unit 24 and the MC 26 is used. Thus, a selected stage of the state machine used in the memory interface 70 of the interface unit 24a is selected. An identical stage of a state machine of one of the interface unit 24b is also selected. The two selected stages are communicated between the interface units 24a, 24b and received by a compare circuit contained in both interface units 24a, 24b. As the interface units operate lock-step with one another, the state machines will likewise march through the same identical states, assuming each state at substantially the same moments in time. If an interface unit encounters an error, or fails, that activity will cause the interface units to diverge, and the state machines will assume different states. The time will come when... ..that will bring to the attention of the CPUs 12A (or 12B) that the interface units 24a, 24b of that CPU are no longer in lock-step, and to act accordingly... ..X port, receiving only those message packets transmitted by the router 14A of the sub-processor system 10A (Fig. 1A). The Y port is serviced by the Y interface unit 24b to receive message packets from the router 14B of the companion sub-processor system 10B. However, both interfaces (as well as Mcs 26 and processor 20), as has been indicated, are basically mirror images of one another in that both... ..in both structure and function. For this reason, message packet information, received by one interface unit (e.g., 24a) must be passed for processing also to the companion interface unit (e.g., 24b). Further, since both interface units 24a, 24b will assemble the same message packets for transmission from the X or the Y ports, the message packet being transmitted by the interface unit (e.g., 24b) actually being communicated from the associated port (e.g., the Y port) will also be coupled to the other interface unit (e.g., 24a) for cross-checking for errors. These features are illustrated in Figs. 6... ..receiving portions of the packet receivers 96 (96x, 96y) of the X and Y interface units 24a, 24b are broadly illustrated. As shown, each packet receiver 96x, 96y has a clock... ..receive a corresponding one of the TNet Links 32. The CS FIFOs 102 operate to synchronize the incoming command/data symbols to the local clock of the packet receiver 96,

buffering... 104x, coupled to the MUX 104y of the packet receiver 96y of the Y interface unit 24b by the cross-link connection 36(sub(x)). In similar fashion, information received at the Y port is coupled to the X interface unit 24a by the cross-link connection 36(sub(y)). In this manner, the command/data... packets received at one of the X, Y ports by the corresponding X, Y, interface unit 24a, 24b is passed to the other so that both will process and communicate the same information on to other components of the interface units 24 and/or memory 28.

Continuing with Fig. 6, depending upon which port X, Y... or the other of the CS FIFOs 102x, 102y for communication to the storage and processing logic 110 of the interface unit 24. The information contained in each 9-bit symbol is an 8-bit byte of... the encoding of which is discussed below with respect to Fig. 9. The storage and processing logic 110 will first translate the 9-bit symbols to 8-bit data or command... the outputs of the CS FIFOs 102x, 102y are also coupled to a command decode unit in addition to the MUX 104. The command decode unit operates to recognize command symbols (differentiating them from data symbols in a manner that is... below), decoding them to generate therefrom command signals that are applied to a receiver control unit, a state machine-based element that functions to control packet receiver operations.

As indicated above... at the output of the MUX 104, the receiver control portion of the storage control unit enables CRC check logic 106 to calculate a CRC symbol while the data symbols are... below, CS FIFOs are found not only in the packet receivers 96 of the interface units 24, but also at each receiving port of the routers 14 and the I/O... an even more important part, and perform a unique function, when a pair of sub-processor systems are operating in duplex mode and the two CPUs 12A and 12B of the sub-processor systems 10A, 10B operate in synchronized, lock-step, executing the same instructions at the same time. When operating in this latter... difficult to ensure that the clocking regime of the routers 14A and 14B are exactly synchronized to those of the CPUs 12A and 12B - even when using frequency locked clocking. In... used to transmit symbols to a CPU 12 and the clock used by an interface unit 24 to receive those symbols.

The structure of the CS FIFO 102 is diagrammatically illustrated... i.e., a packet) or IDLE symbols - except during certain situations (e.g., reset, initialization, synchronization and others discussed below). As explained above, each symbol held in the transmit register 120... same symbol leaving the storage queue, allowing each symbol entering the storage queue 126 to settle before it is clocked out and passed to the storage and processing units 110x (and 110y) by the MUX 104x (and 104y). Since the transmit and receive clocks... functioning in duplex mode) operate to transmit symbols with near frequency clocking. Even so, clock synchronization FIFOs are used at these other ports to receive symbols transmitted with near frequency clocking, and the structure of these clock synchronization FIFOs are substantially the same as that used in frequency locked environments, i.e., that... of the storage queue 126 are nine bits wide; in near frequency environments, the clock synchronization FIFOs use symbol locations of the queue 126 that are 10 bits wide, the extra... the faster clock source. To handle this clock drift, the two pointers are effectively re-synchronized periodically.

When the CPUs 12 are paired and operating in duplex mode, all four interface

units 24 operate in lock-step to, among other things, transmit the same data and receive... simplex mode, each independent of the other, clocking need only be near frequency.

The interface unit 24 receives a SYNC CLK signal that is used in combination with a SYNC command symbol to initialize and synchronize the Rcv register 124 to the transmitting router 14. When using either near frequency or... 102X preferably begin from some known state. Incoming symbols are examined by the storage and processing units 110 of the packet receivers 96. The storage and processing units look for, and act upon as appropriate, command symbols. Pertinent here is that when the... receives a SYNC command symbol it will be decoded and detected by the storage and processing unit 110. Detection of the SYNC command symbol by the storage and processing unit 110 causes assertion of a RESET signal. The RESET signal, under synchronous control of the

SYNC CLK signal, is used to reset the input buffers (including the clock synchronization buffers) to predetermined states, and synchronize them to the routers 14.

The synchronization of the CS FIFOs 102 of the interface units 24 those ...one or both routers 14A, 14B is discussed more fully below in the section discussing synchronization.

Packet Transmitter:

Each interface unit 24 is assigned to transmit from and receive at only one of the X or Y ports of the CPU 12. When one of the interface units 24 transmits, the other operates to check the data being transmitted. This is an important... ...shows, in abbreviated form, the packet transmitters 94x, 94y of the X and Y interface units 24a, 24b, respectively. Both packet transmitters are identically constructed, so that discussion of one (packet... ...logic 152 that receives, from the BTE 88 or AVT 90 of the associated interface unit (here, the X interface unit 24a) the data to be transmitted - in doubleword (64-bit) format. The packet assembly logic... ...and Y ports: they are either symbols that make up a message packet in the process of being transmitted, or IDLE symbols, or other command symbols used to perform control functions... ...154, 156. The output of the multiplexer 154 connects to the X port. (The interface unit 24b connects the output of the multiplexer 154 to the Y port.) The multiplexer 156... ...sub(x)) to the checker logic 160 of the packet transmitter 94y (of the interface unit 24b).

A selection (S) input of the multiplexers receives a 1-bit output from an... ...is accessible to the MP 18 via an OLAP (not shown) formed in the interface unit 24, and is written with information that "personalizes," among other things, the interface units 24 Here, the X/Y stage of the configuration register 162 configures the packet transmitter 94x of the X interface unit 24a to communicate the X encoder 150x output to the X port; the output of... ...traffic is present, the operation of the two packet interfaces 94 (and, thereby, the interface units 24 with which they are associated) are continually monitored. Should one of the checkers detect... ...will be asserted, resulting in an internal interrupt being posted for appropriate action by the processors 20.

Message packet traffic operates in the same manner. Assume, for the moment, that the... ...that information, a byte at a time, to the X encoder 150x of both interface units 96, which will translate each byte to encoded 9-bit form. The output of the... ...is checked with that from the packet transmitter 94x. Again, the operation of the interface units 24a, 24b, and the packet transmitters they contain, are inspected for error.

In the same... ...monitored.

Returning for the moment to Fig. 5, if the outgoing message packet is a processor initiated transaction (e.g., a read request), the processors 20 will expect a message packet to be returned in response. Thus, when the BTE... ...will issue a timeout signal to the interrupt logic (Fig. 14A) to thereby notify the processors 20 of the absence of a response to a particular transaction (e.g., a read... ...the access, to name just a few. Also, the area of memory of the memory unit 28 desired to be accessed are identified in the message packets by virtual or I... ...virtual addresses be translated to physical addresses of the memory 28. Finally, interrupts generated by units or elements external to the CPU 12A, are transmitted via message packets to interrupt the processors 20, which are also written to memory 28 when received. All ...this is handled by the interrupt logic and AVT logic 86, 90.

The AVT logic unit 90 utilizes a table (maintained by the processor 20 in memory 28) containing AVT entries for each possible external source permitted access to the memory 28. Each AVT entry identifies a specific source element or unit and the particular page (a page being nominally 4K (4096) bytes), or portion of a... ...expected" memory accesses. Expected memory accesses are those initiated by the CPU 12 (i.e., processors 20) such as a read request for information from an I/O device. These latter memory accesses are handled by a transaction sequence number (TSN) assigned to each processor initiated request. At about the time the read request is generated, the processors 20 will allocate an area of memory for the data expected to be received in... ...and 26b are, in turn, respectively coupled to the memory interfaces 70 of each interface unit 24a, 24b. The 64-bit doublewords are written

to the memory 28 with the upper... ..check bits respectively from the memory interfaces 70 (70a, 70b) of each of the interface units 24a, 24b (Fig. 5).

Referring to Fig. 10, each memory interface 70 receives, from either the bus 82 from the processor interface 60 or the bus 83 from AVT logic 90 (see Fig. 5), of the associated interface unit 24, 64 bits of data to be written to memory. The busses 76 and 83... ..other for cross-checking between them. Thus, for example, the memory interface 70a (of interface unit 24a) will drive the MC 26a with the "upper" 32 bits of the 64 bits... ..are check bits, leaving 40 bits unused.

Access Validation:

As previously indicated, components of the processing system 10 external to the CPU 12A (e.g., devices of the I/O packet... ..not without qualification. Access validation, as implemented by the AVT logic 90 of the interface units 24, operates to prevent the content of the memory 28 from being ...Accesses to the memory 28 are validated by the AVT logic 90 of each interface unit 24 (Fig. 5), using all of six checks: (1) that the CRC of the message... ..also are permitted the particular message packet source.

The access validation mechanism of the interface unit 24a, AVT logic 88, is shown in greater detail in Fig. 11. Incoming message packets...and post an interrupt to the interrupt logic 86 (Fig. 5) for action by the processor 20.

The mask operation permits the size of the table of AVT entries to be varied. The content of the AVT mask register 175 is accessible to the processor 20, permitting the processors 20 to optionally select the size of the AVT entry table. A maximum AVT table... ..172 allows the AVT size to be matched to the needs of the system. A processing system 10 that includes a larger number of external elements (e.g., the number of... ..amount of the memory space of memory 28 to the AVT entries. Conversely, a smaller processing system 10, with a smaller number of external elements will not have such a large... ..set to a logic "ZERO" indicate an nonexistent TNet address, outside the limits of the processing system 10. A received packet with a TNet address outside the allowable TNet range will... ..in Fig. 11 as being held in the AVT entry register 180 during the validation process. AVT entries have two basic formats: normal and interrupt. The format of a normal AVT...of the AVT input register 170) will result in an error being posted to the processor via an interrupt.

A 12-bit "Permissions" field is included in t AVT entry to... ..path =0). Denials are logged as interrupts with the interrupt logic, and reported to the processor 20 - if the E field is set to a state ("ONE") that enables error-reporting... ..e.g., to a "ONE"), the other fields (Upper Bound, etc.) gain new definitions for processing interrupt writes and managing interrupt queues. This is discussed in more detail below in connection... ..memory 28 will be handled. Set to one state, the requested write operation will be processed normally; set to a second state, write requests specifying addresses with a fractional cache line...be written to a specific queue (interrupt queue) in memory 28, with signalling provided the processors 20 to indicate that an interrupt has been received and "posted," and ready for servicing by the processors 20. Since the interrupt queues are at specific memory locations, the processor can obtain the interrupt data when needed.

An AVT interrupt entry for an interrupt may... ..by the interrupt logic 86, and extracted from the head of the queue by the processor 20 when servicing the interrupt.

The AVT interrupt entry also includes a 20-bit segment ("Source ID") containing source ID information, identifying the external unit seeking attention by the interrupt process. If the source ID information of the AVT interrupt entry does not match that contained... ..class" of the interrupt that is used to determine the interrupt level set in the processor 20 (described more fully below); (2) a queue number that is used to select, as...capability to deliver interrupts to a CPU 12 for servicing. For example, an I/O unit may be unable to complete a read or write transaction issued by a CPU because... ..identify the recipient. These and other errors, exceptions, and irregularities, noted by the I/O units, or the I/O Interface elements, can become the a condition that requires the intervention... ..the AVT

entry register 180 for use by the interrupt logic 86 of the interface unit 24 (Fig. 5), illustrated in greater detail in Fig. 14A.

It is interrupt logic 86...four circular queues specified by the base address information contained in the AVT entry. The processor(s) 20 will then be notified, and it will be up to them as to... ..selected tail queue register 256 by combiner circuit 270, the output of which is the processed by the "mod z" circuit 273 to turn new offset into the queue at which... ..signal. The Queue Full warning signal becomes an "intrinsic" interrupt that is conveyed to the processor units 20 as a warning that if the matter is not promptly handled, later-received interrupt... ..will be discarded.

Incoming message packet interrupts will cause interrupts to be posted to the processor 20 by first setting one of a number of bit positions of an interrupt register 280. Multi-entry queued interrupts are set in interrupt registers 280a for posting to the processor 20; single-entry queue interrupts use interrupt register 280b. Which bit is set depends upon... ..multi-entry queued interrupts, soon after a multi-entry queued interrupt is determined, the interface unit will assert a corresponding interrupt signal (I1) that is applied to decode circuit 283. Decode... ..of register 280a to set, thereby providing advance information concerning the received interrupt to the processor(s) 20, i.e., (1) the type of interrupt posted, and (2) the class of... ..to one another by a compare circuit 279. The update register is writable by the processor 20 to select a register pair for comparison. If the content of the two selected... ..cleared.

Digressing for the moment, there are two basic types of interrupts that concern the processors 20: those interrupts that are communicated to the CPU 12 by message packets, and those...the seven interrupt postings to a latch 288, from which they are coupled to the processor 20 (20a,20b) which has an interrupt register for receiving holding the postings.

In addition... ..change in interrupts (either an interrupt has been serviced, and its posting deleted by the processor 20, or a new interrupt has been posted), a "CHANGE" signal will be issued to the processor interface 60 to inform it that an interrupt posting change has occurred, and that it should communicate the change to the processor 20.

Preferably, the AVT entry register 180 is configured to operate like a single line... ..such as set-associative, fully-associate, or direct-mapped, to name a few.

Coherency:

Data processing systems that use cache memory have long recognized the problem of coherency: making sure that... ..the incoming packet is permitted access are applied to a boundary crossing (Bdry Xing) check unit 219. Boundary check unit 219 also receives an indication of the size of the cache block the CPU 12... ..Len field of the header information from the AVT input register 170. The Bdry Xing unit determines if the data of the incoming packet is not aligned on a cache boundary...time an interrupt will be written to the queued interrupt register 280, to alert the processors 20 that a portion of the incoming data is located in the special queue.

In... ..not, the packet (both header and data) is written to a special queue, and the processors so notified by the intrinsic interrupt process described above. The processors may then move the data from the special queue to cache 22, and later write... ..the cache 22 and the memory 28 is preserved.

Block Transfer Engine (BTE):

Since the processor 20 is inhibited from directly communicating (i.e., sending) information to elements external to the... ..indirect method of information transmission.

The BTE 88 is the mechanism used to implement all processor initiated I/O traffic to transfer blocks of information. The BTE 88 allows creation of... ..BTE registers 300, 302 whose content is coupled to the MUX 306 (of the interface unit 24a; Fig. 5) and used to access the system memory 28 via the memory controllers... ..BTE data

structure 304 in the memory 28 of the CPU 12A (Fig. 2). The processors 20 will write a data structure 304 to the memory 28 each time information is... ..begin on a quadword boundary, and the BTE registers 300, 302 are writable by the processors 20 only. When a processor does write one of the BTE registers 300, 302, it does so with a word... ..the request bit (rc0, rcl) to a clear state, which operates to initiate the BTE process, which is controlled by the BTE state machine 307.

The BTE registers 300, 302 also... ..cause (ec) bit differentiates time-outs and NAKs.

When information is being transferred by the processors 20 to an external unit, the data buffer portion 304b of the data structure 304 holds the information to be transferred. When information from an external unit is received by the processors 20, the data buffer portion 304b is the location targeted to hold the read response information.

The beginning of the data structure 304, portion 304a written by the processor 20, includes an information field (Dest), identifying the external element which will receive the packet... ..the transmitted data is to be written. This information is used by the packet transmitter unit 120 (Fig. 5) to assemble the packet in the form shown in Figs. 3-4...list (el) bit, when set, indicates the end of the chain, and halts the BTE processing.

The interrupt completion (ic) bit, when set, will cause the interface unit 24a to assert an interrupt (BTEComp) which sets a bit in the interrupt register 280... ..the chain pointer).

The interrupt time-out (it) bit, when set, will cause the interface unit 24a to assert an interrupt signal for the processor 20 if the acknowledgement of the access times-out (i.e., if the request timer... ..time), or elicits a NAK response (indicating that the target of the request could not process the request).

Finally, if the check sum (cs) bit is set, the data to be... ..containing the data from which the check sum was formed.

To sum up, when the processors 20 of the CPU 12A desire to send data to an external unit, they will write a data structure 304 to the memory 28, comprising identifier information in portion 304a of the data structure, and the data in the buffer portion 304b. The processors 20 will then determine the priority of the data and will write the BTE register... ..information, and sent.

If the data structure 304 indicates a read request (i.e., the processors 20 are seeking data from an external unit - either an I/O device or a CPU 12), the Len and Local Buffer Ptr... ..receiver 100 (Fig. 5) until the local memory write operation is executed.

Responses to a processor-generated read request to an external unit are not processed by the AVT table logic 146. Rather, when the processors 20 set up the BTE data structure, a transaction sequence number (TSN) is assigned the... ..the BTE 88, which will be an HAC type packet (Fig. 4) discussed above. The processors 20 will also include an memory address in the BTE data structure at which the...302, assume that the foregoing transfer of data from the CPU 12A to an external unit is of a large block of information. Accordingly, a number of data structures would be set up in memory 28 by the processors 20, each (except the last) including a chain pointer to additional data structures, the sum... ..sent. Assume now that a higher priority request is desired to be made by the processors 20. In such a case, the associated data structure 304 for such higher priority request... ..with another BTE operation descriptor.

Memory Controller:

Returning, for the moment, to Fig. 2, interface units 24a, 24b access the memory 28 via a pair of memory controllers (MC) 26a, 26b. The Mcs provide a fail-fast interface between the interface units 24 and the memory 28. The Mcs 26 provide the control logic necessary for accessing... ..in dynamic random access memory (DRAM) logic). The Mcs receive memory requests from the interface units 24, and execute reads and writes as well as providing refresh signals to the DRAMs... ..to provide a 72 bit data path between the memory array 28 and the interface units 24a, 24b, which utilize an SBC-DBD-SbD ECC scheme, where b=4, on a... ..26a, 26b to work together and simultaneously supply a 64-bit word to the interface units 24 with minimum latency, one-half of which (DO) comes

from the MC 26a, and the other half (D1) comes from the other MC 26b. The interface unit 24 generate and check the ECC check bits. The ECC scheme used will not only... 26 bus 25, as well as in internal registers.

From the viewpoint of the interface units 24, the memory 28 is accessed with two instructions: a "read N doubleword" and a... doubleword read or a block read format. The signal called "data valid" tells the interface units 24 two cycles ahead of time that read data is being returned or not being returned.

As indicated above, the maintenance processor (MP 18; Fig. 1A) has two means of access to the CPUs 12. One is... 18 will write a register contained in the OLAP 285 with instructions that permit the processors 20 to build an image of a sequence of instructions in the memory that will permit them (the processors 20) to ...to transfer instructions and data from an external (storage) device that will complete the boot process.

The OLAP 285 is also used by the processors 20 to communicate to the MP 18 error indications. For example, if one of the interface units 24 detect a parity error in data received from the memory controller 26, it will... and address transfers on the bus 25 between the MC 26a and the corresponding interface unit 24a. The addressing and data transfers on the DRAM data bus, as well as generation... the CPU 12.

Packet Routing:

The message packets communicated between the various elements of the processing system 10 (e.g., CPUs 12A, 12B, and devices coupled to the I/O packet... First, each TNet Link L connects to an element (e.g., router 14A) of the processing system 10 via a port that has both receive and transmit capability. Each transmit port... cycle (i.e., each clock period) of the T(underscore)Clk so that the clock

synchronization FIFO at the receiving end of the transmission will maintain synchronization.

Clock synchronization is dependent upon the mode in which the processing system 10 is operated. If operating in the simplex mode in which the CPUs 12A... connect directly to the CPUs may drift with respect to each other. Conversely, when the processing system 10 operates in a duplex mode (e.g., the CPUs operate in synchronized, lock-step operation), the clocks between routers 14 and the CPUs 12 to which they... not necessarily phase-locked).

The flow of data packets between the various elements of the processing system 10 is controlled by command symbols, which may appear at any time, even within... initiated by a CPU 12, or MP 18, and promulgated to all elements of the processing system 10 by the routers 14 to communicate an event requiring software action by all... command symbol is used in conjunction with near frequency operation as an aid to maintaining synchronization between the two clock signals that (1) transfer each symbol to, and load it in each receiving clock synchronization FIFO, and (2) that retrieves symbols from the FIFO.

SLEEP: This command symbol is sent by any element of the processing system 10 to indicate that no additional packet (after the one currently being transmitted, if... received.

SOFT RESET (SRST): The SRST command symbol is used as a trigger during the processes ("synchronization" and "reintegration," described below) that are used to synchronize symbol transfers between the CPUs 12 and the routers 14A, 14B, and then to place... SYNC command symbol is sent by a router 14 to the CPU 12 of the processing system 10 (i.e., the sub-processor systems 10A/10B) to establish frequency-lock synchronization between CPUs 12 and routers 14A, 14B prior to entering duplex mode, or when in duplex mode to request synchronization, as will be discussed more fully below. The SYNC command symbol is used in conjunction... or duplex to simplex), among other things, as discussed further below in the section on Synchronization and Reintegration.

THIS LINK BAD (TLB): When any system element receiving a symbol from a TNet link L (e.g., a router, a CPU, or an I/O unit) notes an error when receiving a command symbol or packet, it will send a TLB... identical pairs of

symbols that are compared to one another when pulled from the clock synchronization FIFOs.. The DVRG command symbol signals the CPU 12 that a mis-compare has been noted. When received by the CPUs, a divergence detection process is entered whereby a determination is made by the CPUs which CPU may be failing... ..command symbols described above operate to control message flow between the various elements of the processing system 10 (e.g., CPUs 12, router 14, and the like), using principally the BUSY however, an "end node" (i.e., a CPU 12 or I/O unit 17 - Fig. 1) may not assert backpressure because one of its transmit ports is backpressured... ..Improperly addressed packets are discarded by the router 14.

When a system element of the processing system 10 receives a BUSY command symbol on a TNet link L on which it... ..other command symbols (READY, BUSY, etc.).

Whenever a TNet port of an element of the processing system 10 detects receipt of a READY command symbol, it will terminate transmission of FILL... ..receives.

As will be seen, all elements (e.g., router 14, CPUs 12) of the processing system 10 that connect to a TNet link L for receiving transmitted symbols will receive those symbols via a clock synchronization (CS) FIFO. For example, as discussed above, the interface units 24 of CPUs 12 include all CS FIFOs 102x, 102y (illustrated in Fig. 6). The... ..depth to allow for speed matching, and the elastic FIFOs must provide sufficient depth for processing delays that may occur between transmission of a BUSY command symbol during receipt of a...another data byte in packet B. As packet A progresses to the next router, the process would be repeated. If the router 14 displaces more data bytes than the FIFO can... ..irrespective of its own findings.

SLEEP Protocol:

The SLEEP protocol is initiated by a maintenance processor via a maintenance interface (an on-line access port - OLAP), described below. The SLEEP protocol... ..reintegrate a slice of the system 10. Routers 14 must be idle (no packets in process) in order to change modes without causing data loss or corruption. When a SLEEP command symbol is received, the receiving element of processing system 10 inhibits initiation of transmission of any new packet on the associated transmit port... ..The HALT command symbol provides a mechanism for quickly informing all CPUs 12 in a processing system 10 that is necessary to terminate HO activity (i.e., message transmissions between the... ..CPUs that receive HALT command symbols on either of their receive ports (of the interface units 24) will post an interrupt to the interrupt register 280 if the system halt interrupt... ..interrupt; Fig. 14A).

The CPUs 12 may be provided with the ability to disable HALT processing. Thus, for example, the configuration registers 75 of the interface units 24 can include a "halt enable register" that, when set to a predetermined state (eg., ZERO) disables HALT processing, but reporting detection of a HALT symbol as an error.

Router Architecture:

Referring now to... ..simplified block diagram of the router 14A is illustrated. The other routers 14 of the processing system 10 (e.g., routers 14B, 14', etc.) are of substantially identical construction and, therefore...these ports 4, 5 are structured to operate in a frequency locked environment when a processing system 10 is set for duplex mode operation. In addition, when in duplex mode, a... ..5)) will receive the command/data symbols from the CPUs, pass them through the clock synchronization FIFOs 518 (discussed further below), and compare each symbol exiting the clock synchronization FIFOs with a gated compare circuit 517. When duplex operation is entered, a configuration register... ..517 to activate the symbol by symbol comparison of the symbols emanating from the two synchronization FIFOs 518 of the router input logic 502 for the ports 4 and 5. Of... ..to that received, at substantially the same time, by the other port input.

To maintain synchronization in the duplex mode, the two port outputs of the router 14A that transmit to... ..mode, are duplicated by the routers 14, and returned to both CPUs.) The output logic units 504(sub(4)), 504(sub(5)) that are coupled directly to the CPUs 12 will... ..message packet identifies only one of the duplexed CPUs 12, e.g., CPU 12A) in synchronized fashion, presenting those symbols in substantially simultaneous fashion to the two CPUs 12.

Of course, the CPUs 12 (more accurately, the associated interface units 24) receive the transmitted symbols with synchronizing FIFOs of substantially the same structure as that illustrated in Fig. 7A so that, even... from the FIFO structures by both CPUs 12 on the same instruction cycle, maintaining the synchronized, lock-step operation of the CPUs 12 required by the duplex operating mode.

As will... conjunction with configuration data written to registers contained in control logic 509 by the maintenance processor 18 (via the on-line access port 285' and serial bus 19A; see Fig. 1A...links L. The input logic 505 of each port input 502 also assists in maintaining synchronization - at least for those ports sending symbols in the near-frequency environment - by removing received... slower-receiving element receiving symbols from a faster-sending element could overload the input clock synchronization FIFO of the slower-receiving element. That is, if a slower clock is used to pull symbols from the clock synchronization FIFO put there by a faster clock, ultimately the clock synchronization FIFO will overflow.

The preferred technique employed here is to periodically insert SKIP symbols in... stream to avoid, or at least minimize, the possibility of an overflow of the clock synchronization FIFO (i.e., clock synchronization FIFO 518; Fig. 20A) of a router 14 (or CPU 12) due to a T... being slightly higher in frequency than the local clock used to pull symbols from the synchronization FIFO. Using SKIP symbols to by-pass a push (onto the FIFO) operation has the... stall each time a SKIP command symbol is received so that, insofar as the clock synchronization FIFO is concerned, the transmitting clock that accompanied the SKIP symbol was missing.

Thus, logic... the port inputs 502 will recognize, and key off receipt of, SKIP command symbols for synchronization in the near frequency clocking environment so that nothing is pushed onto the FIFO, but... 14, or between routers 14, or between a router 14 and an I/O interface unit 16A - Fig. 1) at a 50 Mhz rate, this allows for a worst case frequency... symbol by supplying FILL or IDLE symbols (which are received and pushed onto the clock synchronization FIFOs, but are not passed to the elastic FIFOs). In short, each elastic FIFO 506...received symbols are then communicated from the input register 516 and applied to a clock synchronization FIFO 518, also by the T(underscore)Clk. The clock synchronization FIFO 518 is logically the same as that illustrated in Figs. 8A and 8B, used in the interface units 24 of the CPUs 12. Here, as Fig. 20A shows, the clock synchronization FIFO 518 comprises a plurality of registers 520 that receive, in parallel, the output of... 516. Associated with each of the registers 520 is a two-stage validity (V) bit synchronizer 522, shown in greater detail in Fig. 20B, and discussed below. The content of each... registers 520, together with the one-bit content of each associated two-stage validity bit synchronizer 522, are applied to a multiplexer 524, and the selected register/synchronizer pulled from the FIFO, and coupled to the elastic FIFO 506 by a pair of... is determined the state of the Push Select signal provided by a push pointer logic

unit 530; and, selection of which register 520 will supply its content, via the MUX 524... and loading of the register 520 selected by the push pointer logic 530. Similarly, the synchronization FIFO control logic 534 receives the clock signal local to the router (Rcv Clk) to... pointer logic 532.

Digressing for a moment, and referring to Fig. 20B, the validity bit synchronizer 522 is shown in greater detail as including a D-type flip-flop 541 with... 530 (Fig. 20A) selects the register 520 of the FIFO with which the validity bit synchronizer is associated for receipt of the next symbol - if not a SKIP symbol.

The delay... Truth Table, below). The D-type flip-flop 543 acts as an additional stage of synchronization, ensuring a stable level at the V output relative to the local Rec Clk. The... flip-flop 542, allowing the Pull signal (a periodic pulse from the sync FIFO Control unit 534) to clear the validity bit on this validity synchronizer 522 when the associated register 520 has been read. (Table omitted)

In summary, the validity synchronizer 522 operates to assert a "valid" (V) signal when a symbol is loaded in a...blocked from being routed out a particular port because another message is already in the process of being routed

out that port. However, that other message in turn is also blocked...an incoming message packet bound for the CPUs will be replicated by the crossbar logic unit by routing the message packet to both port output 504(sub(4)) and 504(sub... ..P) identifies which of path (X or Y) should be used for accessing two sub-processing the device.

The routers 14 provide a capability of constructing a large, versatile routing network for, for example, massively parallel processing architectures. Routers are configured according to their location (i.e., level) in the network by...j)) and 509(sub(k)) are such that bits "def" are used in the algorithmic process, then bits "abc" of the Region ID are compared to the content of the Device... ..the route to default register 509(sub(f))) to the final stage of the selection process: check logic 602. Check logic 602 operates to check the status of the port output...a lower level router, and may be located in one or another of the sub-processing systems 10A, 10B. Whether a router is an upper level or lower level router depends... ..of CPUs 12 and I/O devices 16 to one another, forming a massively parallel processing (MPP) system. Other such MPP systems may exist, and it is those routers configured as... ..captured. As soon as the message packet's Destination ID is so captured, the selection process begins, proceeding to the development of a target port address that will be used to...an error that will be posted to the MP18 via the router's (or interface unit's) OLAP for action.

Digressing, it should be appreciated that these protocol rules observed by the routers 14 are also observed by the CPUs 12 (i.e., interface units 24) and I/O packet interfaces 17.

Finally, when the router 14A is in the... ..directly with the CPUs 12A, 12B, and duplex mode is used, a duplex operation logic unit 638 is utilized to coordinate the port output connected to one of the CPUs 12A... ..was able to write instructions to the OLAP 285 that would be executed by the processors 20 to build a small memory image and routine to permit the CPU 12 to... ..the clock generation circuit design. There will be one clock generator circuit in each sub-processor system 10A/10B (Fig. 1) to maintain synchronism. Designated generally with the reference numeral 650... ..used by the various elements (e.g. CPU. 12, routers 14, etc.) of the sub-processor system containing the clock circuit 650 (e.g., 10A).

The clock generator 654 is shown...The 50 Mhz clock signals produced by the counter 663 are distributed throughout the sub-processor system where needed.

Turning now to Fig. 25, there is illustrated the interconnection and use... ..the clock circuits 650 used to develop synchronous clock signals for a pair of sub-processor systems 10A, 10B (Fig. 1) for frequency locked operation. As illustrated in Fig. 25, the two CPUs 12A and 12B of the sub-processor systems 10A, 10B each have a clock circuit 650, shown in Fig. 25 as clock... ..654B of both CPUs 12. A driver and signal line 667 interconnects the two sub-processor systems to deliver the M(underscore)CLK signal developed by the oscillator circuit 652A to the clock generator 654B of the sub-processor system 10B. For fault isolation, and to maintain signal quality, the M(underscore)CLK signal is delivered to the clock generator 654A of the sub-processor system 10A through a separate driver and a loopback connection 668. The reason for the... ..the cable (not shown) will establish the connection shown in Fig. 25 between the sub-processor systems 10A, 10B; connected another way, the connections will be similar, but the oscillator 652B... ..Fig. 25, the M(underscore)CLK signal produced by the oscillator circuit 652A of sub-processing system 10A is used by both sub-processing systems 10A, 10B as their respective SYNC CLK signals and the various other clock signals... ..produced by the clock generators 654A, 654B. Thereby, the clock signals of the paired sub-processing systems 10A, 10B are synchronized for the frequency locked operation necessary for duplex mode.

The VCXOs 662 of the clock... ..This allows both clock generators 654A, 654B to continue to provide to the two sub-processing systems 10A, 10B clock signals in the face of improper operation of the oscillator circuit 652A, although the sub-processor systems may no longer be frequency-locked.

The LOCK signals asserted by the phase comparators... ..LOCK signal signifies that the 50 Mhz signals produced by a clock generator 654 are synchronized, both in phase and in frequency, to the M(underscore)CLK signal. Thus, if either... ..signal that accompanies the symbol stream, and is used to push symbols onto the clock synchronizing

FIFO of the receiving element (router 14, or CPU 12) is substantially identical in frequency... ..not phase, to that of the receiving element used to pull symbols from the clock synchronization FIFOs. For example, referring to Fig. 23, which illustrates symbols being sent from the router... ..clock (Local Clk). The former (Rcv Clk) is used to push symbols onto the clock synchronization FIFOs 126 of each CPU, whereas the latter is used to pull symbols from the... ..much higher frequency clock signal. In such situations provision must be made to ensure that synchronization is maintained between the two CPUs as to symbols pulled from the clock synchronization FIFOs 126 of each.

Here, a constant ratio clocking mechanism is used to control operation of the two clock synchronization FIFOs 126, providing the clock signal that pulls symbols from the two FIFOs at the... ..control mechanism is shown, designated with the reference numeral 70. As Fig. 26A illustrates, clock synchronization FIFO control mechanism 700 includes an pre-settable, multi-stage serial shift register 702, the... ..ratio of the clock signal at which symbols are communicated and pushed onto the clock synchronization FIFOs 126 to the frequency of the clock signal used locally. Here, a 15 stages... ..that will be used as the Local Clk signal to pull symbols from the clock synchronization FIFOs 126, and to operate (update) the pull pointer counter 130. The selected output is... ..of the CPU 12 to the clock signal used to push symbols onto the clock synchronization FIFO 126, Rcv Clk, the serial shift register is preset so that M stages of... ..duplexed CPUs 12 with a 50 Mhz clock. Thus, symbols are pushed onto the clock synchronization FIFOs 126 of the CPUs at a 50 Mhz rate. Assume further that the clock... ..of the MUX 704, which produces the clock signal that pulls symbols from the clock synchronization FIFOs 126, Rcv Clk, will contain, for each 100 ns period, five clock pulses. Thus... ..five symbols will be pushed onto, and five symbols will be pulled from, the clock synchronization FIFOs 126.

This example is symbolically shown in Fig. 26B, while the timing diagram shown... ..labelled "IN" in Fig. 27) of the Rcv Clk will push symbols onto the clock synchronization FIFOs 126. During that same 100 ns period, the serial shift register 702 circulates a... ..clocks which would require additional storage (i.e., an increase in the size of the synchronization FIFO) and impose more latency.

The constant ratio clock circuit presented here (Figs. 26) is... ..frequency to a clock regime of a different, higher frequency. The use of a clock synchronization FIFO is necessary here for compensating effects of signal delays when operating in synchronized, duplexed mode to receive pairs of identical command/data symbols from two different sources. However...so long as there are at least two registers in the place of the clock synchronization FIFO. Transferring data from a higher-frequency clock regime to a lower frequency clock regime... ..a wide range of possible clock ratios.

I/O Packet Interface:

Each of the sub-processor systems 10A, 10B, etc. will have some input/output capability, implemented with various peripheral units, although it is conceivable that the I/O of other sub-processor systems would be available so that a sub-processing system may not necessarily have local I/O. In any event, if local I/O... ..device (e.g., a signal line) would be received by the I/O packet interface unit 16 and used to form an interrupt packet that is sent to the CPU 12... ..OLAP bus, configuration information.

On-Line Access Port:

The MP 18 connects to the interface unit 24, memory controller (MC) 26, routers 14, and I/O packet interfaces with interface signals... ..OLAP 258 is essentially the same, regardless of what element (e.g. router 14, interface unit 24, etc.) it is used with. Fig. 28 diagrammatically illustrates the general structure of the... ..circuit chip used to implement certain of the elements discussed herein. For example, each interface

unit 24, memory controller 26, and router 14 is implemented by an application specific integrated circuit... of the OLAP 158 shown in Fig. 28 describes the OLAP associated with the interface unit 24, the MC 26, and the router 14 of the system.

As Fig. 28 shows...asymmetric variables, a "soft-vote" (SV) logic element 900 (Fig. 30A) is provided each interface unit 24 of each CPU 12. As Fig. 30 illustrates, the SV logic elements 900 of each interface unit 24 are connected to one another by a 2-bit SV bus 902, comprising bus lines 902a and 902b. Bus lines 902a carry one-bit values from the interface units 24 of CPU 12A to those of CPU 12B. Conversely, bus line 902b carries one... the CPU 12A.

Illustrated in Fig. 30B, is the SV logic element 900a of interface unit 24a of CPU 12A. Each SV logic element 900 is substantially identical in construction and... 900a should be understood as applying equally to the other logic elements 900a (of interface unit 24b, CPU 12A), and 900b (of the interface units 24a, 24b of CPU 12B) unless noted otherwise. As Fig. 30B illustrates, the SV logic... the logic elements 900a (as well as its own). In this manner the two interface units 24a, 24b of the CPU 12A can communicate asymmetrical variables to each other.

In a... to the remote register 907 of logic element 902a (and that of the other interface unit 24b).

The logic elements 902 form a part of the configuration registers 74 (Fig. 5). Thus, they may be written by the processor unit(s) 20 by communicating the necessary data/address information over at least a portion of... local and remote registers 906 and 907.

The MUX 914 operates to provide each interface unit 24 of CPU 12A with selective use of the bus line 902a for the SV logic elements 900a, or for communicating a BUS ERROR signal if encountered during the reintegration process (described below) used to bring a pair of CPUs 12 into lock-step, duplex operation... same time, write the enable registers 912 of the logic element 900 of both interface units 24 of each CPU. One of the two logic elements 900 of each CPU will... it is the output enable registers 912 associated with the logic elements 900 of interface units 24a of both CPUs 12A, 12B that are written to enable the associated drivers 916. Thus, the output registers 904 of the interface units 24a of each CPU will be communicated to the bus lines 902; that is, the... to the bus line 902a, while the output register associated with logic element 900b, interface unit 24a of CPU 12B is communicated to bus line 902b. The CPUs 12 will both... again written by each CPU, followed again by reading the remote input registers 907. This process is repeated, one bit at a time, until the entire variable is communicated from the... each CPU 12 to the remote input register of the other. Note that both interface units 24 of CPU 12B will receive the bit of asymmetric information.

One example of use... elements 900 are also used to communicate bus errors that may occur during the reintegration process to be described. When reintegration is being conducted, a REINT signal will be asserted. As... ERROR signal is selected by the MUX 914 and communicated to the bus line 902a.

Synchronization:

Proper operation of the sub-processing systems 10A, 10B (Figs. 1A, 2) whether operating independently (simplex mode), or paired and operating in synchronized lock-step (duplex mode), requires assurance that data communicated between the CPUs 12A, 12B and... routers 14A, 14B will be received properly, and that any initial content of the clock synchronization FIFOs 102 (of CPUs 12A, 12B; Fig. 5) and 519 (of routers 14A, 14B; Fig... erroneously interpreted as data or commands. The push and pull pointers of the various clock synchronization FIFOs 102 (in the CPUs 12) and 518 (in the routers 14) need to be... apart, and presetting the associated FIFO queues to some known state. This done, all clock synchronization FIFOs are initialized for near frequency operation. ...in order to properly implement the lock-step operation of duplex mode operation, the clock synchronization FIFOs must be synchronized to operate with the particular source from which they receive data in order accommodate any... 14A, 14B to the CPUs 12A, 12B must be accounted for. It is the clock synchronization FIFOs 102 of the paired CPUs 12 that operate to receive message packet symbols, adjust... and present symbols to

the two CPUs in a simultaneous manner to maintain lock-step synchronization necessary for duplex mode operation.

In similar fashion, each symbol received by the routers 14A... ..the CPUs (which is discussed further hereinafter). Again, it is the function of the clock synchronization FIFOs 518 of the routers 14A, 14B that receive message packets from the CPUs 12... ..so that the symbols received from the two CPUs 12 are retrieved from the clock synchronization FIFOs simultaneously.

Before discussing how the clock synchronization FIFOs of the CPUs and routers are reset, initialized, and synchronized, an understanding of their operation to maintain synchronous lock-step duplex mode operation is believed helpful. Thus, referring for the moment to Fig. 23, the clock synchronization FIFOs 102 of the CPUs 12A, 12B that receive data, for example, from the router... ..underscore)Clk, from the router 14A to the CPU 12B.

Consider operation of the clock synchronization FIFOs 102(sub(x)), 102(sub(y)), to receive identical symbol streams during duplex operation... ..held by the push and pull pointer counters 128, 130 for the CPU 12A (interface unit 24a), and the content of each of the four storage locations (byte 0. ..., byte 3... ..6 show the same thing for the FIFO 102(sub(y)) of CPU 12B interface unit 24a for each symbol of the duplicated symbol stream.

Assuming the delay 640 is no...0" locations of the queues 126. This is because (1) the FIFOs 102 have been synchronized to operate in synchronism (a process described below), and (2) the push pointer counters 128 are clocked by the clock signal... ..of the symbol stream transmitted by the router 14A will be pulled from the clock synchronization FIFOs 102 of the CPUs 12A, 12B simultaneously, maintaining the required synchronization of received data when operating in duplex mode. In effect, the depths of the queues... ..order to achieve the operation just described with reference to Table 6, the reset and synchronization process shown in Fig 31A is used. The process not only initializes the clock synchronization FIFOs 102 of the CPUs 12A, 12B for duplex mode operation, but also operates to adjust the clock synchronization FIFOs 518 (Fig. 19A) of the CPU ports of each of the routers 14A, 14B for duplex operation. The reset and synchronization process uses the SYNC command symbol to initiate a time period, delineated by the SYNC CLK signal 970 (Fig. 31B), to reset and initialize the respective clock synchronization FIFOs of the CPUs 12A and 12B and routers 14A, 14B. (The SYNC CLK signal... ..It is of a lower frequency than that used to receive symbols by the clock synchronization FIFOs, T(underscore)Clk. For example, where T(underscore)Clk is approximately 50 MHz, the... ..signal is approximately 3.125 MHz.)

Turning now to Fig. 31A, the reset and initialization process begins at step 950 by switching the clock signals used by the CPUs 12A, 12B and routers 14A, 14B as the transmit (T(underscore)Clk) and the unit's local clock (Local Clk) clock signals so that they are derived from the same... ..In addition, configuration registers in the CPUs 12A, 12B (configuration registers 74 in the interface units 24) and the routers 14A, 14B (contained in control logic unit 509 of routers 14A, 14B) are set to the FreqLock state.

The following discussion involves step 952, and makes reference to the interface unit 24 (Fig.5), router 14A (Fig. 19A) and Figs. 31A and 31B. With the clock... ..otherwise be sent followed by a self-addressed message packet. Any message packet in the process of being received and retransmitted when the SLEEP command symbols are received and recognized by... ..per the destination address). The SLEEP command symbol operates to "quiece" router 14A for the synchronization process. The self-addressed message packet sent by the CPU 12A, when received back by the... ..message packet sent after the SLEEP command symbol would necessarily have to be the last processed by the router 14A.

At step 954 the CPU 12A checks to see if it...the router will assert a RESET signal 972 that is applied to the two clock synchronization FIFOs 518 contained in the input logic 505(sub(4)), 505(sub(5)) of the... ..receive symbols directly from CPUs 12A, 12B. RESET, while asserted, will hold the two clock synchronization FIFOs 518 in a temporarily non-operating reset state with the push and pull pointer... ..As each of the CPUs 12 receive SYNC symbols are detected by the storage and processing units of the packet receivers 96 (Figs. 5 and 6) cause the RESET signal to be asserted by the packet receivers 96 (actually, storage and processing elements 110; Fig. 6) of each CPU

12. the RESET signal is applied to the... ..4))), CPUs 12 and routers 14A, 14B de-assert the RESET signals, and the clock synchronization FIFOs of the CPUs 12A, 12, and routers 14A, 14B are released from their reset... ..the delay, the router 14A and CPUs 12 resume pulling data from their respective clock synchronization FIFOs and resume normal operation. The clock synchronization FIFOs of the router 14A begin pulling symbols from the queue (previously set by RESET... ..from the CPU 12A with the T(underscore)Clk will be pushed onto the clock synchronization FIFO at, for example, queue location 0 (or whatever other location pointed to by the... ..0 (or whatever other location the push pointer was set to by RESET). The clock synchronization FIFOs of the router 14A are now synchronized to accommodate whatever delay 640 may be present in one communications path, relative to the... ..and the CPUs 12A, 12B.

Similarly, at the same virtual time, operation of the clock synchronization FIFOs 102 of both CPUs 12A, 12B is resumed, synchronizing them to the router 14A. Also, the CPUs 12A, 12B quit sending the SLEEP command... ..in favor of READY symbols, and resume message packet transmission, as appropriate.

That completes the synchronization process for the router 14A. However, the process must also be performed for the router 14B. Thus, the CPU 12A returns to step... ..however, assuming that the CPUs 12A, 12B are operating in duplex mode, the method and apparatus used to detect and handle a possible error, resulting in divergence of the CPUs from...via a message packet destined for a peripheral device of one or the other sub-processor systems 10A, 10B. Depending upon the destination of the outgoing message packet, step 1002 will... ..router 14 will issue an ERROR signal to the router control logic 509, causing the process to move to step 1004 where the router 14 detecting divergence will transmit a DVRG... ..time outs to occur. A router detecting divergence (without also detecting any simple link error) buys itself time to check the CRC of the received message packet by waiting for the...router 14, or received, all further message packets received from the CPUs and in the process of being routed when divergence was detected, or the DVRG symbol received, will be passed... ..1010) contained in a one of the configuration registers 74 (Fig. 5) of the interface unit 24 of each CPU.

Returning for the moment to step 1006, the determination of which... ..local" is meant to refer to the router 14A, 14B contained in the same sub-processor system 10A, 10B as the CPU. For example, referring to Fig. 1A, router 14A is... ..bit mentioned above: the bit contained in one of the configuration registers 74 of interface unit 24 (Fig. 5) of each CPU. When set to a first state, that particular CPU...the other CPU. In response, the state machines (not shown) within the control and status unit 509 (Fig. 19A) changes the "favorite" bits described above.

A few examples may facilitate understanding... ..DVRG symbol will echo that symbol to the routers 14A, 14B, start its internal divergence process timer, and begin determination of whether to continue or terminate. Having received a TLB symbol...to diverge with no errors reported. This can happen only if software (running on the processors 20) uses known divergent data to alter state. For example, suppose each CPU 12 has... ..number of the CPU 12A will differ from that of the CPU 12B. If the processors use the serial number to change the sequence of instructions executed (say, by branching if the serial number comes after some value) or to modify the value contained in a processor register, the complete "state" of the CPUs 12 will differ. In such cases, the "asymmetrical... ..of the primary CPU simply allows one CPU, and thereby the system 10, to continue processing without software intervention.

- An error at the output of the interface unit 24 of a CPU 12 will be detected by the router 14A, 14B, depending upon... ..router 14A, 14B that connects to a CPU 12 will be detected by the interface unit 24 of the affected CPU. The CPU will send a TLB symbol to the faulty... ..possible failure and, without external intervention, and transparently to the system user, remove the failing unit (CPU 12A or 12B, or router 14A or 14B) from the system to obviate or... ..reintegration." The discussion will refer to the CPUs 12A, 12B, routers 14A, 14B, and maintenance processor 18A, 18B shown forming parts of the processing system 10 illustrated in Fig. 1A. In addition, discussion

will refer to the processors 20a, 20b, the interface units 24a, 24b, and the memory controllers 26a, 26b (Fig. 2) of the CPUs 12A, 12B as single units, since that is the way they function.

Reintegration is used to place two CPUs in...both of the paired CPUs at virtually the same time.

The major steps in the process for changing from simplex mode operation of the one on-line CPU to duplex mode...
...greater detail by the flow diagrams of Figs. 33A - 33D, generally are:

1. Setup and synchronize the two CPUs (one on-line, the other off-line) and their connected routers to... ..the memory of the on-line CPU to the off-line CPU, maintaining a tracking process that monitors changes in the memory of the on-line CPU that have not been... ..and may need to be copied over to, the off-line CPU;
3. Setup and synchronize the CPUs to run a delayed (slave) duplex mode from the same instruction stream (lock... ..will write the predetermined registers (not shown) of the control registers 74 in the interface units 24 of CPUs 12A and 12B, to a next state (after a soft operation) in... ..the off-line CPU 12B.

Next, a sequence is entered (steps 1060 - 1070) that will synchronize the clock synchronization FIFOs of the CPUs 12A, 12B and routers 14A, 14B in much the same fashion... ..the same steps described above in connection with the discussion of Figs. 31A, 31B to synchronize the clock synchronization FIFOs. The on-line CPU 12A will send the sequence of a SLEEP symbol, self-addressed message packet, and SYNC symbol which, with the SYNC CLK signal, operates to synchronize CPUs and routers. Once so synchronized, the on-line CPU 12A then, at step 1066, sends a Soft Reset (SRST) command... ..of all configuration registers and control registers (e.g., configuration registers 74 of the interface units 24) cache, and the like to memory 28 of the on-line ...time to have the system 10 off-line for reintegration. For that reason, the reintegration process is performed in a manner that allows the on-line CPU to continue executing user... ..not match that of the off-line CPU. The reason for this is that normal processing by the processor 20 of the on-line CPU can change memory content after it has been copied... ..when a memory location is written in the on-line CPU 12A during the reintegration process it is marked as "dirty;" second, all copying of memory to the off-line CPU... ..may, however, limit the ability to detect two-bit errors. But, since the memory copying process will last for a only relatively short period of time, this risk is believed acceptable... ..memory location in CPU 12A is made (either an incoming I/O write, or a processor write operation). The returning data (that was copied over to the off-line CPU) would... ..controller 26 (Fig. 2) of the on-line CPU to monitor memory locations in the process of being copied over to the off-line CPU 12B. The memory controller uses a...within the block had been written by another operation (e.g., a write by the processor 20, an I/O write, etc.), that prior write operation will flag the location in... ..still must be copied over to the off-line CPU 12B.

Returning to the reintegration process, and now to Fig. 33B, the memory tracking (AtomicWrite mechanism and using ECC to mark... ..entails writing a reintegration register (not shown; one of the configuration registers 74 of interface unit 24 - Fig. 5) to cause a reintegration (REINT) signal to be asserted. The REINT signal... ..is left alone. Throughout the incremental copy operations, the normal actions of the on-line processor will mark some memory locations dirty.

Several passes of incremental copying will need to be... ..the number of successful WriteConditional operations at the end of each pass through memory, the processors 20 can determine the effect of a given pass compared to the previous pass. When the benefits drop off, the processors 20 will give up on the precopy operations. At this point the reintegration process is ready to place the two CPUs 12A, 12B into lock-step operation.

Thus, the... ..in Fig. 33C, where at step 1100, the on-line CPU 12A momentarily halts foreground processing, i.e., execution of a user application. The remaining state (e.g., configuration registers, cache, etc.) of the on-line processors 20 and its caches is then read and written to a buffer (series of memory... ..to the off-line CPU 12B, together with a "reset vector" that will direct the processor units 20 of both CPUs 12A, 12B to a reset instruction.

Next, step 1106 will quiesce... ..to ensure that the FIFOs of the routers are clear, that the FIFOs of the processor interfaces 24 are clear, and no further incoming I/O message packets are forthcoming. At... ..symbol will be received

and acted upon by both CPUs 12A, 12B, to cause the processor units 20 of each CPU to jump to the location in memory 28 containing the reset... ..a subroutine that will restore the stored state of both CPUs 12A, 12B to the processor units 20, caches 22, registers, etc. The CPUs 12A, 12B will then begin executing the same... ..enabling of the ECC bit to mark dirty locations must now be disabled, since the processors are doing the same thing to the same memory. During this stage of the reintegration... ..encountered by CPU 12A.

Meanwhile, the bus error in the CPU 12A will cause the processor unit 20 to be forced into an error-handling routine to determine (1) the cause of... ..error was caused by an attempt to read a memory location marked dirty. Accordingly, the processor unit 20 will initiate (via the BTE 88 -- Fig. 5) the AtomicWrite mechanism to copy the...the SRST symbols are now received by the CPUs 12A, 12B, they will cause both processor units 20 of the CPUs to be reset to start from the same location with the... ..will periodically update, e.g., a database or audit file that is indicative of the processing of the primary CPU up to that point in time of the update. Should the... ..in error-checking redundancy to the CPU 12B, in the same manner that the individual processor units 20a, 20b of the CPU 12A provide fail-fast, fault tolerance for the CPU - when... ..cost system is applicable, as illustrated in Fig. 34. As shown in Fig. 34, a processing system 10' includes the CPU 12A and routers 14A, 14B structured as described above. The... ..and the CPUs are also the same.

Thus, the CPU 12B' comprises only a single processor unit 20' and associated support components, including the cache 22', interface unit (IU) 24', memory controller 26', and memory 28'. Thus, while the CPU 12A is structured in the manner shown in Fig. 2, with cache processor unit, interface unit, and memory control redundancies, approximately one-half of those components are needed to implement CPU... ..stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... ..i.e, parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of... ..inputs connected to receive the data output

from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... ..for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the... ..attached, without the assistance of any other CPU 12 in the system. Many prior parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O device becomes... ..rest of the system. Other prior systems provide access to I/O through pairs of processors so that should one of the processors fail, ...if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... ..multiprocessing system access to every peripheral, as done here, operates to extend the "primary"/"backup" process taught in the above-identified U.S. Patent No. 4,228,496. There, a multiple CPU system may have a primary process may running on one CPU, while a backup process resides in the background on another of the CPUs. Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be... ..e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system... CPU can access any I/O without the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time... indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Specification: ...both of the paired CPUs at virtually the same time.

The major steps in the process for changing from simplex mode operation of the one on-line CPU to duplex mode... greater detail by the flow diagrams of Figs. 33A - 33D, generally are:

1. Setup and synchronize the two CPUs (one on-line, the other off-line) and their connected routers to... the memory of the on-line CPU to the off-line CPU, maintaining a tracking process that monitors changes in the memory of the on-line CPU that have not been need to be copied over to, the off-line CPU;
3. Setup and synchronize the CPUs to run a delayed (slave) duplex mode from the same instruction stream (lock... will write the predetermined registers (not shown) of the control registers 74 in the interface units 24 of CPUs 12A and 12B, to a next state (after a soft operation) in... the off-line CPU 12B.

Next, a sequence is entered (steps 1060 - 1070) that will synchronize the clock synchronization FIFOs of the CPUs 12A, 12B and routers 14A, 14B in much the same fashion... the same steps described above in connection with the discussion of Figs. 31A, 31B to synchronize the clock synchronization FIFOs. The on-line CPU 12A will send the sequence of a SLEEP symbol, self-addressed message packet, and SYNC symbol which, with the SYNC CLK signal, operates to synchronize CPUs and routers. Once so synchronized, the on-line CPU 12A then, at step 1066, sends a Soft Reset (SRST) command... of all configuration registers and control registers (e.g., configuration registers 74 of the interface units 24) cache, and the like to memory 28 of the on-line CPU, copying the... time to have the system 10 off-line for reintegration. For that reason, the reintegration process is performed in a manner that allows the on-line CPU to continue executing user... not match that of the off-line CPU. The reason for this is that normal processing by the processor 20 of the on-line CPU can change memory content after it has been copied... when a memory location is written in the on-line CPU 12A during the reintegration process it is marked as "dirty;" second, all copying of memory to the off-line CPU... may, however, limit the ability to detect two-bit errors. But, since the memory copying process will last for a only relatively short period of time, this risk is believed acceptable... memory location in CPU 12A is made (either an incoming I/O write, or a processor write operation). The returning data (that was copied over to the off-line CPU) would... controller 26 (Fig. 2) of the on-line CPU to monitor memory locations in the process of being copied over to the off-line CPU 12B. The memory controller uses a... within the block had been written by another operation (e.g., a write by the processor 20, an I/O write, etc.), that prior write operation will flag the location in... still must be copied over to the off-line CPU 12B.

Returning to the reintegration process, and now to Fig. 33B, the memory tracking (AtomicWrite mechanism and using ECC to mark... entails writing a reintegration register (not shown; one of the configuration registers 74 of interface unit 24 - Fig. 5) to cause a reintegration (REINT) signal to be asserted. The REINT signal... is left alone. Throughout the incremental copy operations, the normal actions of the on-line processor will mark some memory locations dirty.

Several passes of incremental copying will need to be... ..the number of successful WriteConditional operations at the end of each pass through memory, the processors 20 can determine the effect of a given pass compared to the previous pass. When the benefits drop off, the processors 20 will give up on the precopy operations. At this point the reintegration process is ready to place the two CPUs 12A, 12B into lock-step operation.

Thus, the... ..in Fig. 33C, where at step 1100, the on-line CPU 12A momentarily halts foreground processing, i.e., execution of a user application. The remaining state (e.g., configuration registers, cache, etc.) of the on-line processors 20 and its caches is then read and written to a buffer (series of memory... ..to the off-line CPU 12B, together with a "reset vector" that will direct the processor units 20 of both CPUs 12A, 12B to a reset instruction.

Next, step 1106 will quiesce... ..to ensure that the FIFOs of the routers are clear, that the FIFOs of the processor interfaces 24 are clear, and no further incoming I/O message packets are forthcoming. At... ..symbol will be received and acted upon by both CPUs 12A, 12B, to cause the processor units 20 of each CPU to jump to the location in memory 28 containing the reset... ..a subroutine that will restore the stored state of both CPUs 12A, 12B to the processor units 20, caches 22, registers, etc. The CPUs 12A, 12B will then begin executing the same... ..enabling of the ECC bit to mark dirty locations must now be disabled, since the processors are doing the same thing to the same memory. During this stage of the reintegration... ..encountered by CPU 12A.

Meanwhile, the bus error in the CPU 12A will cause the processor unit 20 to be forced into an error-handling routine to determine (1) the cause of... ..error was caused by an attempt to read a memory location marked dirty. Accordingly, the processor unit 20 will initiate (via the BTE 88 -- Fig. 5) the AtomicWrite mechanism to copy the... ..the SRST symbols are now received by the CPUs 12A, 12B, they will cause both processor units 20 of the CPUs to be reset to start from the same location with the... ..will periodically update, e.g., a database or audit file that is indicative of the

processing of the primary CPU up to that point in time of the update. Should the... ..in error-checking redundancy to the CPU 12B, in the same manner that the individual processor units 20a, 20b of the CPU 12A provide fail-fast, fault tolerance for the CPU - when... ..cost system is applicable, as illustrated in Fig. 34. As shown in Fig. 34, a processing system 10' includes the CPU 12A and routers 14A, 14B structured as described above. The... ..and the CPUs are also the same.

Thus, the CPU 12B' comprises only a single processor unit 20' and associated support components, including the cache 22', interface unit (IU) 24', memory controller 26', and memory 28'. Thus, while the CPU 12A is structured in the manner shown in Fig. 2, with cache processor unit, interface unit, and memory control redundancies, approximately one-half of those components are needed to implement CPU... ..stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... ..i.e., parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of... ..inputs connected to receive the data output from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... ..for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the... ..attached, without the assistance of any other CPU 12 in the system. Many prior parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O device becomes... ..rest of the system. Other prior systems provide access to I/O through pairs of processors so that should one of the processors fail, access to the corresponding I/O is still available through the remaining I/O... ..if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... ..multiprocessing system access to every peripheral, as done here, operates to extend the "primary"/"backup" process taught in the above-identified U.S. Patent No. 4,228,496. There, a multiple CPU system may have a primary process running on one CPU, while a backup process resides in the background on another of the CPUs. Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be... ..e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system... ..CPU can access any I/O without the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... ..four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... ..an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time...indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... ..When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... ..response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Claims: ...A2

1. In a processing system comprising at least a pair of transmitting/receiving elements coupled to one another by...

6/K/28 (Item 28 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Patent Assignee:

- Compaq Computer Corporation... ;
;;

Country	Number	Kind	Date	
Type	Pub. Date		Kind	Text
...Compaq Computer Corporation (687797) 20555 SH 249 Houston,		19		

Texas 77070-2698 US			
Application...	19		
Available Text	Language	Update	Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: ...and 08/485,053, filed concurrently herewith.

The present invention is directed generally to data processing systems, and more particularly to a multiple processing system and a reliable system area network that provides connectivity for interprocessor and input/output... ..and communications systems to general purpose high availability commercial systems. The evolution of fault tolerant computers has been well documented (see D. P. Siewiorek, R. S. Swarz, "The Theory and Practice... ..and the Jet Propulsion laboratory began to apply fault tolerance to the development of guidance computers for aerospace applications. The 1960's also saw the development of the first AT&T electronic switching systems.

The first commercial fault tolerant machines were introduced by Tandem Computers in the 1970's for use in on-line transaction processing applications (J. Bartlett, "A NonStop Kernal," in proc. Eighth Symposium on Operating System Principles, pp... ..systems were introduced in the 1980's (O. Serlin, "Fault- Tolerant Systems in Commercial Applications," Computer, pp. 19-30, August 1984). Current commercial fault tolerant systems include distributed memory multi-processors, shared-memory transaction based systems, "pair-and-spare" hardware fault tolerant systems (see R. Freiburghouse, "Making Processing Fail-safe," Mini-micro Systems, pp. 255-264, May 1982; U.S. Patent No. 4... ..system.), and triple-modular-redundant systems such as the "Integrity" computing system manufactured by Tandem Computers Incorporated of Cupertino, California, assignee of this application and the invention disclosed herein.

Most applications of commercial fault tolerant computers fall into the category of on-line transaction processing. Financial institutions require high availability for electronic funds transfer, control of automatic teller machines, and... ..telecommunications systems.

Vendors of fault tolerant machines attempt to achieve both increased system availability, continuous processing, and correctness of data even in the presence of faults. Depending upon the particular system architecture, application software ("processes") running on the system either continue to run despite failures, or the processes are automatically restarted from a recent checkpoint when a fault is encountered. Some fault tolerant systems are provided with sufficient component redundancy to be able reconfigure around failed components, but processes running in the failed modules are lost. Vendors of commercial fault tolerant systems have extended fault tolerance beyond the processors and disks. To make large improvements in reliability, all sources of failure must be addressed... ..power supplies, fans and inter-module connections.

The "NonStop," and "Integrity" architectures manufactured by Tandem Computers Incorporated, (both respectively illustrated broadly in U.S. Patent No. 4,228,496 and U... ..assigned to the assignee of this application; NonStop and Integrity are registered trademarks of Tandem Computers Incorporated) represent two current approaches to commercial fault tolerant computing. The NonStop system, as generally... ..above-identified U.S. Patent No. 4,278,496, employs an architecture that uses multiple processor systems designed to continue operation despite the failure of any single hardware component. In normal operation, each processor system uses its major components independently and concurrently, rather than as "hot backups". The NonStop system architecture may consist of up to 16 processor systems interconnected by a bus for interprocessor communication. Each processor system has its own memory which contains a copy of a message-based operating system. Each processor system controls one or more input/output (I/O) busses. Dual-porting of I/O controllers and devices provides multiple paths to each device.

External storage (to the processor system), such as disk storage, may be mirrored to maintain redundant permanent data storage.

This... hardware, while fault recovery is the responsibility of the software.

Also, in the Nonstop multi-processor architecture, application software ("process") may run on the system under the operating system as "process-pairs," including a primary process and a backup process. The primary process runs on one of the multiple processors while the backup process runs on a different processor. The backup process is usually dormant, but periodically updates its state in response to checkpoint messages from the primary process. The content of a checkpoint message can take the form of complete state update, or currently most application code runs under transaction processing software which provides recovery through a combination of checkpoints and transaction two-phase commit protocols.

Interprocessor message traffic in the Tandem Nonstop architecture includes each processor periodically broadcasting an "I'm Alive" message for receipt by all the processors of the system, including itself, informing the other processors that the broadcasting processor is still functioning. When a processor fails, that failure will be announced and identified by the absence of the failed processor's periodic "I'm Alive" message. In response, the operating system will direct the appropriate backup processes to begin primary execution from the last checkpoint. New backup processes may be started in another processor, or the process may be run with no backup until the hardware has been repaired. U.S. Patent... example of this technique.

Each I/O controller is managed by one of the two processors to which it is attached. Management of the controller is periodically switched between the processors. If the managing processor fails, ownership of the controller is automatically switched to the other processor. If the controller fails, access to the data is maintained through another controller.

In addition to providing hardware fault tolerance, the processor pairs of the above-described architecture provide some measure of software fault tolerance. When a processor fails due to a software error, the backup processor frequently is able to successfully continue processing without encountering the same error. The software environment in the backup processor typically has different queue lengths, table sizes, and process mixes. Since most of the software bugs escaping the software quality assurance tests involve infrequent data dependent boundary conditions, the backup processes often succeed.

In contrast to the above-described architecture, the Integrity system illustrates another approach... fault recovery is the logical choice since few modifications to the software are required. The processors and local memories are configured using triple-modular-redundancy (TMR). All processors run the same code stream, but clocking of each module is independent to provide tolerance... three streams is asynchronous, and may drift several clock periods apart. The streams are re-synchronized periodically and during access of global memory. Voters on the TMR Controller boards detect and mask failures in a processor module. Memory is partitioned between the local memory on the triplicated processor boards and the global memory on the duplicated TMRC boards. The duplicated portions of the... techniques to detect failures. Each global memory is dual ported and is interfaced to the processors as well to the I/O Processors (IOPs). Standard VME peripheral controllers are interfaced to a pair of busses through a Bus... the BIMs to switch control of all controllers to the remaining IOP. Mirrored disk storage units may be attached to two different VME controllers.

In the Integrity system all hardware failures... reintegrated on-line.

The preceding examples illustrate present approaches to incorporating fault tolerance into data processing systems. Approaches involving software recovery require less redundant hardware, and offer the potential for some... have been developed on other systems.

Thus, the systems described above provide fault tolerant data processing either by hardware (e.g, fail-functional, employing redundancy) or by software techniques (fail-fast... hardware). However, none of the systems described

are believed capable of providing fault tolerant data processing, using both hardware (fail-functional) and software (fail-fast) approaches, by a single data processing system.

Computing systems, such as those described above, are often used for electronic commerce: electronic data interchange (EDI) and global messaging. Today's demands upon such electronic commerce, however, is demanding more and more throughput capacity as the number of users increases and... networks such as local area networks (LANS), and the like.

A key requirement for a server architecture is the ability to move massive quantities of data. The server should have high bandwidth that is scalable, so that added throughput capacity can be added... response time, latency affects service levels and employee productivity.

The present invention provides a multiple-processor system that combines both of the two above-described approaches to fault tolerant architecture, hardware redundancy and software recovery techniques, in a single system.

Broadly, the present invention includes a processing system composed of multiple sub-processing systems. Each sub-processing system has, as the main processing element, a central processing unit (CPU) that in turn comprises a pair of processors operating in lock-step, synchronized fashion... execute each instruction of an instruction stream at the same time. Each of the sub-processing systems further include an input/output (I/O) system area network system that provides redundant communication paths between various components of the larger processing

system, including a CPU and assorted peripheral devices (e.g., mass storage

units, printers, and the like) of a sub-processing system, as well as between the sub-processors that may make up the larger overall processing system. Communication between any component of the processing system (e.g., a CPU and another CPU, or a CPU and any peripheral device, regardless of which sub-processing system it may belong to) is implemented by forming and transmitting packetized messages that are... responsible for choosing the proper or available communication paths from a transmitting component of the processing system to a destination component based upon information contained in the message packet. Thus, the... peripherals, but permits it to also be used for interprocessor communications.

As indicated above, the processing system of the present invention is structured to provide fault-tolerant operation through both "fail... at a variety of points in the various data paths between the (lock-step operated) processor elements of the CPU and its associated memory. In particular, the processing system of the present invention conducts error-checking at an interface, and in a manner... little impact on performance. Prior art systems typically implement error-checking by running pairs of processors, and checking (comparing) the data and instruction flow between the processors and a cache memory. This technique of error-checking tended to add delay to the... error-checking precluded use of off-the-shelf parts that may be available (i.e., processor/cache memory combinations on a single semiconductor chip or module). The present invention performs error-checking of the processors at points that operate at slower rates, such as the main memory and I/O interfaces which operate at slower speeds than the processor-cache interface. In addition, the error-checking is performed at locations that allow detection of errors that may occur in the processors, their cache memory, and the I/O and memory interfaces. This allows simpler designs for... other data integrity checks.

Error-checking of the communication flow between the components of the processing system is achieved by adding a cyclic-redundancy-check (CRC) to the message packets that... Good" (TPG) or "This Packet Bad" (TPB) - is appended to every packet. A maintenance diagnostic processor can use this information to isolate a link or router element that introduces an error... of topologies, so that alternate paths can be provided between any two elements of a processing system (e.g., between a CPU and an I/O device), for communication in the... so (e.g., by creating a "deadlock" condition, discussed further below).

The CPUs of a processing system are capable of operating in one of two basic modes: a "simplex mode" in...
...independently of the other, or a "duplex "mode in which pairs of CPUs operate in synchronized, lock-step fashion. Simplex mode operation provides the capability of recovering from faults that are... U.S. Pat. No. 4,228,496 which teaches a multiprocessing system in which each processor has the capability of checking on the operability of its sibling processors, and of taking over the processing of a processor found or believed to have failed). When operating in duplex mode, the paired CPUs both... fault tolerant platform for less robust operating systems (e.g., the UNIX operating system). The processing system of the present invention, with the paired, lock-step CPUs, is structured so that... masked (i.e., operating despite the existence of a fault), primarily through hardware.

When the processing system is operating in duplex mode, each CPU pair uses the I/O system to access any peripheral of the processing system, regardless of which (of the two, or more) sub-processor system the peripheral may be ostensibly a member of. Also, in duplex mode, message packets... message for the CPU pair (from either a peripheral device such as a mass storage unit or from a processing unit), will replicate the message and deliver it to both CPUs of the pair using synchronization methods that ensure that the CPUs remain synchronized. In effect, the duplex CPU pair, as viewed from the I/O system and other... as a single CPU. Thus, the I/O system, which includes elements from all sub-processing systems, is made to be seen by the duplex CPU pair as one homogeneous system... a multiprocessor system in which the CPU of any one is actually a pair of synchronized, lock-step CPUs.

Yet another important aspect of the present invention is that interrupts issuing... interrupts via the message packet system ensures that they will arrive at duplexed CPUs in synchronized fashion, in the same manner as I/O message packets. Interrupt message packets will contain... the system. In addition, using the same messaging system to communicate data between I/O units and the CPUs and to communicate interrupts to the CPUs preserves the ordering of I... the implementation of a technique of validating access to the memory of any CPU. The processing system, as structured according to the present invention, permits the memory of any CPU to a CPU and any other component of the processor system. Thereby, the individual processor units of the CPU are removed from the more mundane tasks of getting information from memory and out onto the TNet network, or accepting information from the network. The processor unit of the CPU merely sets up data structures in memory containing the data to be... is required, where in memory the response is to be placed when received. When the processor unit completes the task of creating the data structure, the block transfer engine is notified to... response is received, it is routed to the expected memory location identified, and notifies the processor unit that the response was received.

Further aspects and features of the present invention will become... invention, which should be taken in conjunction with the accompanying drawings.

Fig. 1A illustrates a processing system constructed in accordance with the teachings of the present invention, and Figs. 1B and 1C illustrate two alternate configurations of the processing system of Fig. 1A, employing clusters or arrangements of the processing system of Fig. 1A;

Fig. 2 illustrates, in simplified block diagram form, the central processing unit (CPU) that forms a part of each sub-processor system of Figs. 1A - 1C;

Figs. 3A - 3D and 4A - 4C illustrate the construction of... the area network I/O system shown in Fig. 2;

Fig. 5 illustrates the interface unit that forms a part of the CPUs of Fig. 2 to interface the processor and memory with the I/O area network system;

Fig. 6 is a block diagram, illustrating a portion of packet receiver of the interface unit of Fig. 5;

Fig. 7A diagrammatically illustrates the clock synchronization FIFO (CS FIFO) used by the packet receiver section packet receiver shown in Fig. 6;

Fig. 7B is an block diagram of a construction of the clock synchronization FIFO structure shown in Fig. 7A;

Fig. 8 illustrates the cross-connections for error-checking outbound transmissions from the two interface units of a CPU;

Fig. 9 illustrates an encoded (8B to 9B) data/command symbol;

Fig. 10 illustrates the method and structure used by the interface unit of Fig. 5 to cross-check for errors data being transferred to the memory controllers... of a CPU of Fig. 2 to other (external to the CPU) components of the processing system;

Fig. 12 is a block diagram that diagrammatically illustrates the formation of an address... 14A illustrates the logic for posting interrupt requests to queues in memory and to the processor units of the CPU of Fig. 2;

Fig. 14B illustrates the process used to form a memory address for a queue entry;

Fig. 15 is a block... data output constructs formed in the memory of the CPU of Fig. 2 by a processor unit, and containing data to be sent via the area I/O networks shown in Figs. 1A - 1C, and also illustrating the block transfer engine (BTE) unit of the interface unit of Fig. 5 that operates to access the data output constructs for transmission to the... pair of memory controllers between memory of a CPU of Fig. 2 and its interface unit for accessing from memory 72 bits of data, including two simultaneously-accessed 32-bit words... other for error-checking;

Fig. 19A is a simplified block diagram illustration of the router unit used in the area input/output networks of the processing systems shown in Figs. 1A - 1C;

Fig. 19B illustrates comparison on two port inputs of the router unit of Fig. 19A;

Fig. 20A is a block diagram the construction of one of the six input ports of the router unit shown in Fig. 19A;

Fig. 20B is a block diagram of the synchronization logic used to validate command/data symbols received at an input port of the router unit of Fig. 19A;

Fig. 21A is a block diagram illustration of the target port selection... is a block diagram illustration of one of the six output ports of the router unit shown in Fig. 19A;

Fig. 23 is an illustration of the method used to transmit identical information to a duplexed pair CPUs of Fig. 2 in synchronized fashion when the processing system is operating in lock-step (duplex) mode, using a pair the FIFOs of Fig... is a simplified block diagram illustrating the clock generation system of each of the sub-processing systems of Figs. 1A - 1C for developing the plurality of clock signals used to operate the various elements of that sub-processing system;

Fig. 25 illustrates the topology used to interconnect the clock generation systems of paired sub-processing systems for synchronizing the various clock signals of the pair of sub-processing systems to one another;

Fig. 26A and 26B illustrates a FIFO constant rate clock control logic used to control the clock synchronization FIFO of Figs. 8 or 20 in the situation when the two clocks used to... structure of the on-line access port (OLAP) used to provide access to the maintenance

processor (MP) to the various elements of the system of Fig. 1A (or those of Figs... the soft-flag logic used to handle asymmetric variables between the CPUs of paired sub-processing systems operating in duplex mode;

Fig. 31A shows a flow diagram, and Fig. 31B illustrates a portion of SYNC CLK, both of which are used to reset and synchronize the clock synchronization FIFOs of the CPUs and routers of the processing system of Fig. 1A that receive information from each other;

Fig. 32 is a flow... 33A - 33D generally illustrate the procedure used to bring one of the CPUs of processing system shown in Fig. 1A into lock-step, duplex mode operation with the other of the CPUs without measurably halting operation of the processing system; and

Fig. 34 illustrates a reduced cost architecture incorporating teachings of the invention; and... to the figures and, for the moment, principally Fig. 1A, there is illustrated a data processing system, designated with the reference 10, constructed according to the various teachings of the present invention. As Fig. 1A shows, the data processing system 10 comprises two sub-processor systems 10A and 10B each of which are substantially the same in structure and function... should be appreciated that, unless noted otherwise, a description of any one of the sub-processor systems 10 will apply equally to any other sub-processor system 10.

Continuing with Fig. 1A therefore, each of the sub-processor systems 10A, 10B is illustrated as including a central processing unit (CPU) 12, a router 14, and a plurality of input/output (I/O) packet interfaces... one of the I/O packet interfaces 16 will also have coupled thereto a maintenance processor (MP) 18.

The MP 18 of each sub-processor system 10A, 10B connects to each of the elements of that sub-processor system via an IEEE 1149.1 test bus 17 (shown in phantom in Fig. 1A... accompanying clock signal. As Fig. 1A further illustrates, TNet Links L also interconnect the sub-processor systems 10A and 10B to one another, providing each sub-processor system 10 with access to the I/O devices of the other as well as inter-CPU communication. As will be seen, any CPU 12 of the processing system 10 can be given access to the memory of any other CPU 12, although... the memory of a CPU 12 by a wayward peripheral device 17.

Preferably, the sub-processor systems 10A/10B are paired as illustrated in Fig. 1A (and Figs 1B and 1C, discussed below), and each sub-processor system 10A/10B pair (i.e., comprising a CPU 12, at least one router 14... 12A) connects, by a TNet Link L to a router (14A) of the corresponding sub-processor system (e.g., 10A). Conversely, the Y port connects the CPU (12A) to the router (14B) of the companion sub-processor system (10B). This latter connection not only provides a communication path for access by a CPU (12A) to the I/O devices of the other sub-processor system (10B), but also to the CPU (12B) of that system for inter-CPU communication.

Information is communicated between any element of the processing system 10 and any other element (e.g., CPU 12A of sub-processor system 10A) of the system and any other element of the system (e.g., an I/O device associated with an I/O packet interface 16B of sub-processor system 10B) via message "packets." Each message packet is made up of a number of... this reason, a unique method of receiving the symbols at the receiver, using a clock synchronization first-in-first-out (CS FIFO) storage structure (described more fully below), has been developed... operation means just that: the frequencies of the clock signals of the transmitter and receiver units are locked, although not necessarily in phase. Frequency locked clock signals are used to transmit symbols between the routers 14A, 14B and the CPUs 12 of paired sub-processor systems (e.g., sub-processor systems 10A, 10B, Fig. 1A). Since the clocks of the transmitting and receiving element are not phase related, a clock synchronization FIFO is again used -- albeit operating in a slightly different mode from that used for... difference, as will be seen, is due to the fact that pairs of the sub-processor systems 10 can be operated in a synchronized, lock-step mode, called duplex mode, in which each CPU 12 operates to execute the... 1A illustrates another feature of the invention: a cross-link connection between the two sub-processor systems 10A, 10B through the use of additional routers 14 (identified in Fig. 1A as... RY(sub(1)), and RY(sub(2)) form a cross-link connection between the sub-processors 10A, 10B (or, as shown, "sides" X and Y, respectively) to couple them to I... the routers RX(sub(2)) and RY(sub(2)) provide the I/O packet interface units 16x and 16y with a dual ported interface. Of course, it will now be evident... lend themselves to being used in a manner that can extend the configuration of the processing system 10 to include additional sub-processor systems such as illustrated in Figs. 1B and 1C. In Fig. 1B, for example, one... of each of the routers 14A and 14B is used to connect the corresponding sub-processor systems 10A and 10B to additional sub-processor systems 10A' and 10B' forming thereby a larger processing system comprising clusters of the basic processing system 10 of Fig. 1.

Similarly, in Fig. 1C the above concept is extended to form an eight sub-processor system cluster, comprising sub-processor systems pairs 10A/10B, 10A'/10B', 10A"/10B", and 10A'''/10B'''. In turn, each of the sub-processor systems (e.g., sub-processor system 10A) will have essentially the same basic minimum configuration of a CPU 12, a... ..by a I/O packet interface 16, except that, as Fig. 1C shows, the sub-processor systems 10A and 10B include additional routers 14C and 14D, respectively, in order to extend the cluster beyond sub-processor systems 10A'/10B' to the sub-processor systems 10A"/10B" and 10A'''/10B''' . As Fig. 1C further illustrates, unused ports 4 and... ..the routers 14 when configuring the topology of the system 10, any CPU 12 of processing system 10 of Fig. 1C can access any other "end unit" (e.g., a CPU or I/O device) of any of the other sub-processor systems. Two paths are available from any CPU 12 to the last router 14 connecting to the I/O packet interface 16. For example, the CPU 12B of the sub-processor system 10B' can access the I/O 16''' of sub-processor system 10A''' via router 14B (of sub-processor system 10B'), router 14D, and router 14B (of sub-system 10B''') and, via link LA... ..10A'''), OR via router 14A (of sub-system 10A'), router 14C, and router 14A (sub-processor system 10A'''). Similarly, CPU 12A of sub-processor system 10A" may access (via two paths) memory contained in the CPU 12B of sub-processor 10B to read or write data. (Memory accesses by one CPU 12 of another component of the processing system requires, as will be seen, the components seeking access to have authorization to do... ..prevents corruption of memory data of a CPU by erroneous access.)

The topology of the processing system shown in Fig. 1B is achieved by using port 1 of the routers 14A, 14B, and auxiliary TNet links LA, to connect to the routers 14A', 14B' of sub-processor systems 10A', 10B'. The topology thereby obtained establishes redundant communication paths between any CPU 12 (12A, 12B, 12A', 12B') and any I/O packet interface 16 of the processing system 10 shown in Fig. 1B. For example, the CPU 12A' of the sub-processor system 10A' may access the I/O 16A of sub-processor system 10A by a first path formed by the router 14A' (in port 4, out... ..shown in Fig. 1B. By interconnecting one port of each router 14 of each sub-processor pair, and using additional auxiliary TNet links LA (illustrated in Fig. 1C with the dotted line connections) between the ports 1 of the routers 14 (14A" and 14B") of sub-processor systems 10A", 10B" and 10A''', 10B''', two separate, independent data paths can be found between any CPU 12 and any I/O packet interface 16. In this fashion, any end unit (i.e., a CPU 12 or an I/O packet interface 16) will have at least two paths to any other end unit.

Providing alternate paths of access between any two end units (e.g., between a CPU 12 and any other CPU 12, or between any CPU... ..any two of the remaining fault domains. Here, a fault domain could be a sub-processor system (e.g., 10A). Thus, if the sub-processor system 10A were brought down because of a failure the electrical power being supplied, without... ..TNet link LA between the routers 14A''' and 14B''', the CPU 12B of the sub-processor system 10B would have lost access to the I/O packet interface 16''' (via router... ..with the loss of the router 14A (and router 14C) by loss of the sub-processor system 10A, communications between the CPU 12B is still possible via the route of router... ..equally to CPU 12B. As Fig. 2 shows, the CPU 12A includes a pair of processor units 20a, 20b that are configured for synchronized, lock-step operation in that both processor units 20a, 20b receive and execute identical instructions, and issue identical data and command outputs, at substantially the same moments in time. Each of the processor units 20a and 20b is connected, by a bus 21 (21a, 21b) to a corresponding cache memory 22. The particular type of processor units used could contain sufficient internal cache memory so that the cache memory 22 would not... ..22 could be used to supplement any cache memory that may be internal to the processor units 20. In any event, if the cache memory 22 is used, the bus 21 is... ..22 address bits, 3 bits of parity covering the address, and 7 control bits.

The processors 20a, 20b are also respectively coupled, via a separate 64-bit address/data bus 23 to X and Y interface units 24a, 24b. If desired, the address/data communicated on each bus 23a, 23b could also be protected by parity, although this will increase the width of the bus. (Preferably, the processors 20 are constructed to include RISC R4000 type microprocessors, such as are available from the MIPS Division of Silicon Graphics, Inc. of Santa Clara, California.)

The X and Y interface units 24a, 24b operate to communicate data and command signals between the processor units 20a, 20b and a memory system of the CPU 12A, comprising a memory controller (MC... ..MC halves 26a and 26b) and a dynamic random access memory array 28. The interface units 24 interconnect to each other and to the Mcs 26a, 26b by a 72-bit... ..accompanied by 8 bits of ECC) are written to the memory 28 by the interface units 24, one interface unit 24 will drive only one word (e.g., the 32 most significant portion) of the doubleword being written while the other interface unit 24 writes the other word of the double word (e.g., the least significant 32-bit portion of the doubleword). In addition, on each write operation the interface units 24a, 24b perform a cross-check operation on the data not written by that interface unit 24 with the data written by the other to check for errors; on read operations... ..accessed corresponds to the address of the location from which the doubleword was stored.

Interface units 24a, 24b of the CPU 12A form the circuitry to respectively service the X and Y (I/O) ports of the CPU 12A. Thus, the X interface unit 24a connects by the bi-directional TNet Link Lx to a port of the router 14A of the processor system 10A (Fig. 1A) while the Y interface unit 24b similarly connects to the router 14B of the processor system 10B by TNet Link Ly. The X interface unit 24a handles all I/O traffic between the router 14A and the CPU 12A of the sub-processor system 10A. Likewise, the Y interface unit 24b is responsible for all I/O traffic between the CPU 12A and the router 14B of companion sub-processor system 10B.

The TNet Link Lx connecting the X interface unit 24a to the router 14A (Fig. 1) comprises, as above indicated, two 10-bit buses... ..sub(x)) carries data incoming from the router 14A. In similar fashion, the Y interface unit 24b is connected to the router 14B (of the sub-processor system 10B) by two 10-bit busses: 30(sub(y)) (for outgoing transmissions) and 32... ..y)) (for incoming transmissions), together forming the TNet Link Ly.

The X and Y interface units 24a, 24b are synchronously operated in lock-step, performing substantially the same operations at substantially the same times. Thus, although only the X interface unit 24a actually transmits data onto the bus 30(sub(x)), the same output data is being produced by the Y interface unit 24b, and used for error-checking. The Y interface unit 24b output data is coupled to the X interface unit 24a by a cross-link 34(sub(y)) where it is received by the X interface unit 24a and compared against the same output data produced by the X interface unit. In this way the outgoing data made available at the X port of the CPU... ..the port of the CPU 12A is checked. The output data from the Y interface unit 24b is coupled to the Y port by a 10-bit bus 30(sub(y)), and also to the X interface unit 24a by the 9-bit cross-link 34(sub(y)) where is checked with that produced by the X interface unit.

As mentioned, the two interface units 24a, 24b operate in synchronous, lock-step with one another, each performing substantially the same... ..X and/or Y ports of the CPU 12A must be received by both interface units 24a, 24b to maintain the two interface units in this lock-step mode. Thus, data received by one interface unit 24a, 24b is passed to the other, as indicated by the dotted lines and 9... ..sub(x)) (communicating incoming data being received at the X port by the X interface unit 24a to the Y interface unit 24b) and 36(sub(y)) (communicating data received at the Y port by the Y interface unit 24b to the X interface unit 24a).

Certain more robust operating systems are structured with a fault-tolerant capability in the... ..example, U.S. Patent No. 4,817,091 teaches a multiprocessor system in which each processor periodically messages each of the processors of the system (including itself), under software control, to thereby provide an indication of continuing operation. Each of the processors, in addition to performing its normal tasks, operates as a backup processor to another of the processors. In the event one of the backup processors fails to receive the messaged indication from a sibling processor, it will take over the operation of that sibling (now thought to be inoperative), in... ..platform for both types of software. Thus, when a robust operating system is available, the processing system 10 can be configured to operate in a "simplex" mode in which each of... ..left, in most instances, to software.

Alternatively, for less robust operating systems and software, the processing system 10 provides a hardware-based fault-tolerance by being configured to operate in a... ..g., CPUs 12A, 12B) are coupled together as shown in Fig. 1A, to operate in synchronized, lock-step fashion, executing the same instructions at the substantially the same moment in time... ..data and command symbols. In order to simplify the design of the CPU 12, the processors 20 are

precluded from communicating directly with any outside entity (e.g., another CPU 12... ..O device via the I/O packet interface 16). Rather, as will be seen, the processor will construct a data structure in memory and turn over control to the interface units 24. Each interface unit 24 includes a block transfer engine (BTE; Fig. 5) configured to provide a form of... ..to the destination according to information contained in the message packet.

The design of the processing system 10 permits a memory 28 of a CPU to be read or written by... ..via the routers 14. Accordingly, before continuing with the description of the construction of the processing system 10, it would be of advantage to understand first the configuration of the data...information.

As indicated, the HADC message packet operates to communicate write data between the end units (e.g., CPU 12) of the processing system 10. Other message packets, however, may be differently constructed because of their function and... ..CRC. The HC message packet is used to acknowledge a request to write data.

Interface Unit:

The X and Y interface units 24 (i.e., 24a and 24b - Fig. 2) operate to perform three major functions within the CPU 12: to interface the processors 20 to the memory 28; to provide an I/O service that operates transparently to, but under the control of, the processors; and to validate requests for access to the memory 28 from outside sources.

Regarding first the interface function, the X and Y interface units 24a, 24b operate to respectively communicate processors 20a, 20b to the memory controllers (Mcs 26a, 26b) and memory 28 for writing and... ..fast checking of the data read/written. For example, write operations have the two interface units 24a, 24b cooperating to cross-check the data to be written to ensure its integrity (and at the same time, the interface units 24 will operate) to develop an error correcting code (ECC) that covers, as will be appropriate address.

With respect to I/O access, the processors 20 are not provided with the ability to communicate directly with the input/output systems... ..must write data structures to the memory 28 and then pass control to the interface units 24 which perform a direct memory access (DMA) operation to retrieve those data structures, and... ..indicated in the data structure itself.)

The third function of the X and Y interface units 24, access validation to the memory 28, uses an address validation and translation (AVT) table maintained by the interface units. The AVT table contains an address for each system component (e.g., an I/O... ..the incoming message packets are virtual addresses. These virtual addresses are translated by the interface unit to physical addresses recognizable by the memory control units 26 for accessing the memory 28.

Referring to Fig. 5, illustrated is a simplified block diagram of the X interface unit 24a of the CPU 12A. The companion Y interface unit 24b (as well as the interface units 24 of the CPU 12B, or any other CPU 12) is of substantially identical construction. Accordingly, it will be understood that a description of the interface unit 24a will apply equally to the other interface units 24 of the processing system 10.

As Fig. 5 illustrates, the X interface unit 24a includes a processor interface 60, a memory interface 70, interrupt logic 86, a block transfer engine (BTE) 88, access validation and translation logic 90, a packet transmitter 94, and a packet receiver 96.

Processor Interface:

The processor interface 60 handles the information flow (data and commands) between the processor 20a and the X interface unit 24a. A processor bus 23, including a 64 bit address and data bus (SysAD) 23a and a 9 bit command bus 23b, couples the processor 20a and the processor interface 60 to one another. While the SysAD bus 23a carries memory address and data... ..and qualifying commands carried at substantially the same time on the SysAD bus 23a. The processor interface 60 operates to interpret commands issued by the processor unit 20a in order to pass reads/writes to memory or control registers of the processor interface. In addition, the processor interface 60 contains temporary storage (not shown) for buffering addresses and data for access to... ..26). Data and command

information read from memory is similarly buffered en route to the processor unit 20a, and made available when the processor unit is ready to accept it. Further, the processor interface 60 will operate to generate the necessary interrupt signalling for the X interface unit 24a.

The processor interface 60 is connected to a memory interface 70 and to configuration registers 74 by a bi-directional 64 bit processor address/data bus 76. The configuration registers 74 are a symbolic representation of the various control registers contained in other components of the X interface unit 24a, and will be discussed when those particular components are discussed. However, although not specifically... ..throughout other of the logic that is used to implement the X interface 24a, the

processor address/data bus 76 is likewise coupled to read or write to those registers.

Configuration registers 74 are read/write accessible to the processor 20a; they allow the X interface unit to be "personalized." For example, one register identifies the node address of the CPU 12A... ..with the CPU 12A; another, readable only, contains a fixed identification number of the interface unit 24, and still other registers define areas of memory that can be used by, for... ..logic 90, etc.) employing them are discussed.

The memory interface 70 couples the X interface unit 24a to the memory controllers 26 (and to the Y interface unit 24b; see fig. 2) by a bus 25 that includes two 36 bi-directional bit... ..25a, 25b. The memory interface operates to arbitrate between requests for memory access from the processor unit 20, the BTE 88, and the AVT logic 90. In addition to memory accesses from the processor unit 20a, the memory 28 may also be accessed by components of the processing system 10 to, for example, store data requested to be read by the processor unit 20a from an I/O unit 17, or memory 28 may also be accessed for I/O data structures previously set up in memory by the processor unit. Since these accesses are all asynchronous, they must be arbitrated, and the memory interface 70... ..command information accessed from the memory 28 is coupled from the memory interface to the processor interface 60 by a memory read bus 82, as well as to an interrupt logic... ..doubleword quantities. However, while the memory interfaces 70 of both the X and Y interface units 24a ...by the memory interface 70 are coupled to the memory interface by the companion interface unit 24 where they are compared with the same 32 bits for error.

Digressing for the... ..containing interrupt information are received, that information is conveyed to the interrupt logic 86 for processing and posting for action by the processor 20, along with any interrupts generated internal to the CPU 12A. Internally generated interrupts will... ..register 71 (internal to the interrupt logic 86), indicating the cause of the interrupt. The processor 20 can then read and act upon the interrupt. The interrupt logic is discussed more fully below.

The BTE 88 of the X interface unit 24a operates to perform direct memory accesses, and provides the mechanism that allows the processors 20 to access external resources. The BTE 88 can be set-up by the processors 20 to generate I/O requests, transparent to the processors 20 and notify the processors when the requests are complete. The BTE logic 88 is discussed further below.

Requests for... ..8 byte wide format necessary for storing in the memory 28.

Outgoing message packets containing processor originated transaction requests (e.g., a read request asking for a block data from an I/O unit) are monitored by the request transaction logic (RTL) 100. The RTL 100 provides a time... ..will generate an interrupt (handled and reported by the interrupt logic 86) to inform the processor 20 that the request was not honored. In addition, the RTL 100 will validate responses... ..28 (by the DMA operation of the BTE 86) at a location known to the processor 20 so that it can locate the response.

Each of the CPUs 12 are checked... ..discussed. One such check is an on-going monitor of the operation of the interface units 24a, 24b of each CPU. Since the interface units 24a, 24b operate in lock-step synchronism checking can be performed by monitoring the operating states of the paired interface units 24a, 24b by a continuous

comparison of certain of their internal states. This approach is implemented by using one stage of a state machine (not shown) contained in the unit 24a of CPU 12A, and comparing each state assumed by that stage with its identical state machine stage in the interface unit 24b. All units of the interface units 24 use state machines to control their operations. Preferably, therefore, a state machine of the memory interface 70 that controls the data transfers between the interface unit 24 and the MC 26 is used. Thus, a selected stage of the state machine used in the memory interface 70 of the interface unit 24a is selected. An identical stage of a state machine of one of the interface unit 24b is also selected. The two selected stages are communicated between the interface units 24a, 24b and received by a compare circuit contained in both interface units 24a, 24b. As the interface units operate lock-step with one another, the state machines will likewise march through the same identical states, assuming each state at substantially the same moments in time. If an interface unit encounters an error, or fails, that activity will cause the interface units to diverge, and the state machines will assume different states. The time will come when... ..that will bring to the attention of the CPUs 12A (or 12B) that the interface units 24a, 24b of that CPU are no longer in lock-step, and to act accordingly... ..X port, receiving only those message packets transmitted by the router 14A of the sub-processor system 10A (Fig. 1A). The Y port is serviced by the Y interface unit 24b to receive message packets from the router 14B of the companion sub-processor system 10B. However, both interfaces (as well as Mcs 26 and processor 20), as has been indicated, are basically mirror images of one another in that both... ..in both structure and function. For this reason, message packet information, received by one interface unit (e.g., 24a) must be passed for processing also to the companion interface unit (e.g., 24b). Further, since both interface units 24a, 24b will assemble the same message packets for transmission from the X or the Y ports, the message packet being transmitted by the interface unit (e.g., 24b) actually being communicated from the associated port (e.g., the Y port) will also be coupled to the other interface unit (e.g., 24a) for cross-checking for errors. These features are illustrated in Figs. 6... ..receiving portions of the packet receivers 96 (96x, 96y) of the X and Y interface units 24a, 24b are broadly illustrated. As shown, each packet receiver 96x, 96y has a clock... ..receive a corresponding one of the TNet Links 32. The CS FIFOs 102 operate to synchronize the incoming command/data symbols to the local clock of the packet receiver 96, buffering... ..104x, coupled to the MUX 104y of the packet receiver 96y of the Y interface unit 24b by the cross-link connection 36(sub(x)). In similar fashion, information- received at the Y port is coupled to the X interface unit 24a by the cross-link connection 36(sub(y)). In this manner, the command/data... ..packets received at one of the X, Y ports by the corresponding X, Y, interface unit 24a, 24b is passed to the other so that both will process and communicate the same information on to other components of the interface units 24 and/or memory 28.

Continuing with Fig. 6, depending upon which port X, Y... ..or the other of the CS FIFOs 102x, 102y for communication to the storage and processing logic 110 of the interface unit 24. The information contained in each 9-bit symbol is an 8-bit byte of... ..the encoding of which is discussed below with respect to Fig. 9. The storage and processing logic 110 will first translate the 9-bit symbols to 8-bit data or command... ..the outputs of the CS FIFOs 102x, 102y are also coupled to a command decode unit in addition to the MUX 104. The command decode unit operates to recognize command symbols (differentiating them from data symbols in a manner that is... ..below), decoding them to generate therefrom command signals that are applied to a receiver control unit, a state machine-based element that functions to control packet receiver operations.

As indicated above... ..at the output of the MUX 104, the receiver control portion of the storage control unit enables CRC check logic 106 to calculate a CRC symbol while the data symbols are... ..below, CS FIFOs are found not only in the packet receivers 96 of the interface units 24, but also at each receiving port of the routers 14 and the I/O...an even more important part, and perform a unique function, when a pair of sub-processor systems are operating in duplex mode and the two CPUs 12A and 12B of the sub-processor systems 10A, 10B operate in synchronized, lock-step, executing the same instructions at the same time. When operating in this latter... ..difficult to ensure that the clocking regime of the routers 14A and 14B are exactly synchronized to those of the CPUs 12A and 12B - even when using frequency locked clocking. In... ..used to transmit symbols to a CPU 12 and the clock used by an interface unit 24 to receive those symbols.

The structure of the CS FIFO 102 is diagrammatically illustrated... ..i.e., a packet) or IDLE symbols - except during certain situations (e.g., reset, initialization, synchronization and others discussed below). As explained above, each symbol held in the transmit register 120...same symbol leaving the storage queue, allowing each symbol entering the storage queue 126 to settle before it is clocked out and passed to the storage and processing units 110x (and 110y) by the MUX 104x (and 104y). Since the transmit and receive clocks... ..functioning in duplex mode) operate to transmit symbols with near frequency clocking. Even so, clock synchronization FIFOs are used at these other ports to receive symbols transmitted with near frequency clocking, and the structure of these clock synchronization FIFOs are substantially the same as that used in frequency locked environments, i.e., that... ..of the storage queue 126 are nine bits wide; in near frequency environments, the clock synchronization FIFOs use symbol locations of the queue 126 that are 10 bits wide, the extra... ..the faster clock source. To handle this clock drift, the two pointers are effectively re-synchronized periodically.

When the CPUs 12 are paired and operating in duplex mode, all four interface

units 24 operate in lock-step to, among other things, transmit the same data and receive... ..simplex mode, each independent of the other, clocking need only be near frequency.

The interface unit 24 receives a SYNC CLK signal that is used in combination with a SYNC command symbol to initialize and synchronize the Rcv register 124 to the transmitting router 14. When using either near frequency or... ..102X preferably begin from some known state. Incoming symbols are examined by the storage and processing units 110 of the packet receivers 96. The storage and processing units look for, and act upon as appropriate, command symbols. Pertinent here is that when the... ..receives a SYNC command symbol it will be decoded and detected by the storage and processing unit 110. Detection of the SYNC command symbol by the storage and processing unit 110 causes assertion of a RESET signal. The RESET signal, under synchronous control of the SYNC CLK signal, is used to reset the input buffers (including the clock synchronization buffers) to predetermined states, and synchronize them to the routers 14.

The synchronization of the CS FIFOs 102 of the interface units 24 those ...one or both routers 14A, 14B is discussed more fully below in the section discussing synchronization.

Packet Transmitter:

Each interface unit 24 is assigned to transmit from and receive at only one of the X or Y ports of the CPU 12. When one of the interface units 24 transmits, the other operates to check the data being transmitted. This is an important... ..shows, in abbreviated form, the packet transmitters 94x, 94y of the X and Y interface units 24a, 24b, respectively. Both packet transmitters are identically constructed, so that discussion of one (packet... ..logic 152 that receives, from the BTE 88 or AVT 90 of the associated interface unit (here, the X interface unit 24a) the data to be transmitted - in doubleword (64-bit) format. The packet assembly logic... ..and Y ports: they are either symbols that make up a message packet in the process of being transmitted, or IDLE symbols, or other command symbols used to perform control functions... ..154, 156. The output of the multiplexer 154 connects to the X port. (The interface unit 24b connects the output of the multiplexer 154 to the Y port.) The multiplexer 156... ..sub(x)) to the checker logic 160 of the packet transmitter 94y (of the interface unit 24b).

A selection (S) input of the multiplexers receives a 1-bit output from an... ..is accessible to the MP 18 via an OLAP (not shown) formed in the interface unit 24, and is written with information that "personalizes," among other things, the interface units 24 Here, the X/Y stage of the configuration register 162 configures the packet transmitter 94x of the X interface unit 24a to communicate the X encoder 150x output to the X port; the output of... ..traffic is present, the operation of the two packet interfaces 94 (and, thereby, the interface units 24 with which they are associated) are continually monitored. Should one of the checkers detect... ..will be asserted, resulting in an internal interrupt being posted for appropriate action by the processors 20.

Message packet traffic operates in the same manner. Assume, for the moment, that the... ..that information, a byte at a time, to the X encoder 150x of both interface units 96, which will translate each byte to encoded 9-bit form. The output of the... ..is checked with that from the packet transmitter 94x. Again, the operation of the interface units 24a, 24b, and the packet transmitters they contain, are inspected for error.

In the same... ..monitored.

Returning for the moment to Fig. 5, if the outgoing message packet is a processor initiated transaction (e.g., a read request), the processors 20 will expect a message packet to be returned in response. Thus, when the BTE... ..will issue a timeout signal to the interrupt logic (Fig. 14A) to thereby notify the processors 20 of the absence of a response to a particular transaction (e.g., a read... ..the access, to name just a few. Also, the area of memory of the memory unit 28 desired to be accessed are identified in the message packets by virtual or I... ..virtual addresses be translated to physical addresses of the memory 28. Finally, interrupts generated by units or elements external to the CPU 12A, are transmitted via message packets to interrupt the processors 20, which are also written to memory 28 when received. All ...this is handled by the interrupt logic and AVT logic 86, 90.

The AVT logic unit 90 utilizes a table (maintained by the processor 20 in memory 28) containing AVT entries for each possible external source permitted access to the memory 28. Each AVT entry identifies a specific source element or unit and the particular page (a page being nominally 4K (4096) bytes), or portion of a... ..expected" memory accesses. Expected memory accesses are those initiated by the CPU 12 (i.e., processors 20) such as a read request for information from an I/O device. These latter memory accesses are handled by a transaction sequence number (TSN) assigned to each processor initiated request. At about the time the read request is generated, the processors 20 will allocate an area of memory for the data expected to be received in... ..and 26b are, in turn, respectively coupled to the memory interfaces 70 of each interface unit 24a, 24b. The 64-bit doublewords are written to the memory 28 with the upper... ..check bits respectively from the memory interfaces 70 (70a, 70b) of each of the interface units 24a, 24b (Fig. 5).

Referring to Fig. 10, each memory interface 70 receives, from either the bus 82 from the processor interface 60 or the bus 83 from AVT logic 90 (see Fig. 5), of the associated interface unit 24, 64 bits of data to be written to memory. The busses 76 and 83... ..other for cross-checking between them. Thus, for example, the memory interface 70a (of interface unit 24a) will drive the MC 26a with the "upper" 32 bits of the 64 bits... ..are check bits, leaving 40 bits unused.

Access Validation:

As previously indicated, components of the processing system 10 external to the CPU 12A (e.g., devices of the I/O packet... ..not without qualification. Access validation, as implemented by the AVT logic 90 of the interface units 24, operates to prevent the content of the memory 28 from being ...Accesses to the memory 28 are validated by the AVT logic 90 of each interface unit 24 (Fig. 5), using all of six checks: (1) that the CRC of the message... ..also are permitted the particular message packet source.

The access validation mechanism of the interface unit 24a, AVT logic 88, is shown in greater detail in Fig. 11. Incoming message packets...and post an interrupt to the interrupt logic 86 (Fig. 5) for action by the processor 20.

The mask operation permits the size of the table of AVT entries to be varied. The content of the AVT mask register 175 is accessible to the processor 20, permitting the processors 20 to optionally select the size of the AVT entry table. A maximum AVT table... ..172 allows the AVT size to be matched to the needs of the system. A processing system 10 that includes a larger number of external elements (e.g., the number of... ..amount of the memory space of memory 28 to the AVT entries. Conversely, a smaller processing system 10, with a smaller number of external elements will not have such a large... ..set to a logic "ZERO" indicate a nonexistent TNet address, outside the limits of the processing system 10. A received packet with a TNet address outside the allowable TNet range will... ..in Fig. 11 as being held in the AVT entry register 180 during the validation process. AVT entries have two basic

formats: normal and interrupt. The format of a normal AVT...of the AVT input register 170) will result in an error being posted to the processor via an interrupt.

A 12-bit "Permissions" field is included in the AVT entry to...path =0). Denials are logged as interrupts with the interrupt logic, and reported to the processor 20 - if the E field is set to a state ("ONE") that enables error-reporting...e.g., to a "ONE"), the other fields (Upper Bound, etc.) gain new definitions for processing interrupt writes and managing interrupt queues. This is discussed in more detail below in connection...memory 28 will be handled. Set to one state, the requested write operation will be processed normally; set to a second state, write requests specifying addresses with a fractional cache line...be written to a specific queue (interrupt queue) in memory 28, with signalling provided the processors 20 to indicate that an interrupt has been received and "posted," and ready for servicing by the processors 20. Since the interrupt queues are at specific memory locations, the processor can obtain the interrupt data when needed.

An AVT interrupt entry for an interrupt may...by the interrupt logic 86, and extracted from the head of the queue by the processor 20 when servicing the interrupt.

The AVT interrupt entry also includes a 20-bit segment ("Source ID") containing source ID information, identifying the external unit seeking attention by the interrupt process. If the source ID information of the AVT interrupt entry does not match that contained...class" of the interrupt that is used to determine the interrupt level set in the processor 20 (described more fully below); (2) a queue number that is used to select, as...capability to deliver interrupts to a CPU 12 for servicing. For example, an I/O unit may be unable to complete a read or write transaction issued by a CPU because...identify the recipient. These and other errors, exceptions, and irregularities, noted by the I/O units, or the I/O Interface elements, can become the a condition that requires the intervention...the AVT entry register 180 for use by the interrupt logic 86 of the interface unit 24 (Fig. 5), illustrated in greater detail in Fig. 14A.

It is interrupt logic 86...four circular queues specified by the base address information contained in the AVT entry. The processor(s) 20 will then be notified, and it will be up to them as to...selected tail queue register 256 by combiner circuit 270, the output of which is the processed by the "mod z" circuit 273 to turn new offset into the queue at which...signal. The Queue Full warning signal becomes an "intrinsic" interrupt that is conveyed to the processor units 20 as a warning that if the matter is not promptly handled, later-received interrupt...will be discarded.

Incoming message packet interrupts will cause interrupts to be posted to the processor 20 by first setting one of a number of bit positions of an interrupt register 280. Multi-entry queued interrupts are set in interrupt registers 280a for posting to the processor 20; single-entry queue interrupts use interrupt register 280b. Which bit is set depends upon...multi-entry queued interrupts, soon after a multi-entry queued interrupt is determined, the interface unit will assert a corresponding interrupt signal (I1) that is applied to decode circuit 283. Decode...of register 280a to set, thereby providing advance information concerning the received interrupt to the processor(s) 20, i.e., (1) the type of interrupt posted, and (2) the class of...to one another by a compare circuit 279. The update register is writable by the processor 20 to select a register pair for comparison. If the content of the two selected...cleared.

Digressing for the moment, there are two basic types of interrupts that concern the processors 20: those interrupts that are communicated to the CPU 12 by message packets, and those...the seven interrupt postings to a latch 288, from which they are coupled to the processor 20 (20a,20b) which has an interrupt register for receiving holding the postings.

In addition...change in interrupts (either an interrupt has been serviced, and its posting deleted by the processor 20, or a new interrupt has been posted), a "CHANGE" signal will be issued to the processor interface 60 to inform it that an interrupt posting change has occurred, and that it should communicate the change to the processor 20.

Preferably, the AVT entry register 180 is configured to operate like a single line... such as set-associative, fully-associate, or direct-mapped, to name a few.

Coherency:

Data processing systems that use cache memory have long recognized the problem of coherency: making sure that... the incoming packet is permitted access are applied to a boundary crossing (Bdry Xing) check unit 219. Boundary check unit 219 also receives an indication of the size of the cache block the CPU 12... Len field of the header information from the AVT input register 170. The Bdry Xing unit determines if the data of the incoming packet is not aligned on a cache boundary... time an interrupt will be written to the queued interrupt register 280, to alert the processors 20 that a portion of the incoming data is located in the special queue.

In... not, the packet (both header and data) is written to a special queue, and the processors so notified by the intrinsic interrupt process described above. The processors may then move the data from the special queue to cache 22, and later write... the cache 22 and the memory 28 is preserved.

Block Transfer Engine (BTE):

Since the processor 20 is inhibited from directly communicating (i.e., sending) information to elements external to the... indirect method of information transmission.

The BTE 88 is the mechanism used to implement all processor initiated I/O traffic to transfer blocks of information. The BTE 88 allows creation of... BTE registers 300, 302 whose content is coupled to the MUX 306 (of the interface unit 24a; Fig. 5) and used to access the system memory 28 via the memory controllers... BTE data structure 304 in the memory 28 of the CPU 12A (Fig. 2). The processors 20 will write a data structure 304 to the memory 28 each time information is... begin on a quadword boundary, and the BTE registers 300, 302 are writable by the processors 20 only. When a processor does write one of the BTE registers 300, 302, it does so with a word... the request bit (rc0, rc1) to a clear state, which operates to initiate the BTE process, which is controlled by the BTE state machine 307.

The BTE registers 300, 302 also... cause (ec) bit differentiates time-outs and NAKs.

When information is being transferred by the processors 20 to an external unit, the data buffer portion 304b of the data structure 304 holds the information to be transferred. When information from an external unit is received by the processors 20, the data buffer portion 304b is the location targeted to hold the read response information.

The beginning of the data structure 304, portion 304a written by the processor 20, includes an information field (Dest), identifying the external element which will receive the packet... the transmitted data is to be written. This information is used by the packet transmitter unit 120 (Fig. 5) to assemble the packet in the form shown in Figs. 3-4... list (el) bit, when set, indicates the end of the chain, and halts the BTE processing.

The interrupt completion (ic) bit, when set, will cause the interface unit 24a to assert an interrupt (BTEComp) which sets a bit in the interrupt register 280... the chain pointer).

The interrupt time-out (it) bit, when set, will cause the interface unit 24a to assert an interrupt signal for the processor 20 if the acknowledgement of the access times-out (i.e., if the request timer... time), or elicits a NAK response (indicating that the target of the request could not process the request).

Finally, if the check sum (cs) bit is set, the data to be... containing the data from which the check sum was formed.

To sum up, when the processors 20 of the CPU 12A desire to send data to an external unit, they will write a data structure 304 to the memory 28, comprising identifier information in portion 304a of the data structure, and the data in the buffer portion 304b. The processors 20 will then determine the priority of the data and will write the BTE register... information, and sent.

If the data structure 304 indicates a read request (i.e., the processors 20 are seeking data from an external unit - either an I/O device or a CPU 12), the Len and Local Buffer Ptr... receiver 100 (Fig. 5) until the local memory write operation is executed.

Responses to a processor-generated read request to an external unit are not processed by the AVT table logic 146. Rather, when the processors 20 set up the BTE data structure, a transaction sequence number (TSN) is assigned the... the BTE 88, which will be an HAC type packet (Fig. 4) discussed above. The processors 20 will also include an memory address in the BTE data structure at which the...302, assume that the foregoing transfer of data from the CPU 12A to an external unit is of a large block of information. Accordingly, a number of data structures would be set up in memory 28 by the processors 20, each (except the last) including a chain pointer to additional data structures, the sum... sent. Assume now that a higher priority request is desired to be made by the processors 20. In such a case, the associated data structure 304 for such higher priority request... with another BTE operation descriptor.

Memory Controller:

Returning, for the moment, to Fig. 2, interface units 24a, 24b access the memory 28 via a pair of memory controllers (MC) 26a, 26b. The Mcs provide a fail-fast interface between the interface units 24 and the memory 28. The Mcs 26 provide the control logic necessary for accessing... in dynamic random access memory (DRAM) logic). The Mcs receive memory requests from the interface units 24, and execute reads and writes as well as providing refresh signals to the DRAMs... to provide a 72 bit data path between the memory array 28 and the interface units 24a, 24b, which utilize an SBC-DBD-SbD ECC scheme, where $b=4$, on a... 26a, 26b to work together and simultaneously supply a 64-bit word to the interface units 24 with minimum latency, one-half of which (DO) comes from the MC 26a, and the other half (D1) comes from the other MC 26b. The interface unit 24 generate and check the ECC check bits. The ECC scheme used will not only... 26 bus 25, as well as in internal registers.

From the viewpoint of the interface units 24, the memory 28 is accessed with two instructions: a "read N doubleword" and a... doubleword read or a block read format. The signal called "data valid" tells the interface units 24 two cycles ahead of time that read data is being returned or not being returned.

As indicated above, the maintenance processor (MP 18; Fig. 1A) has two means of access to the CPUs 12. One is... 18 will write a register contained in the OLAP 285 with instructions that permit the processors 20 to build an image of a sequence of instructions in the memory that will permit them (the processors 20) ...to transfer instructions and data from an external (storage) device that will complete the boot process.

The OLAP 285 is also used by the processors 20 to communicate to the MP 18 error indications. For example, if one of the interface units 24 detect a parity error in data received from the memory controller 26, it will... and address transfers on the bus 25 between the MC 26a and the corresponding interface unit 24a. The addressing and data transfers on the DRAM data bus, as well as generation... the CPU 12.

Packet Routing:

The message packets communicated between the various elements of the processing system 10 (e.g., CPUs 12A, 12B, and devices coupled to the I/O packet... First, each TNet Link L connects to an element (e.g., router 14A) of the processing system 10 via a port that has both receive and transmit capability. Each transmit port... clock cycle (i.e., each clock period) of the T Clk so that the clock

synchronization FIFO at the receiving end of the transmission will maintain synchronization.

Clock synchronization is dependent upon the mode in which the processing system 10 is operated. If operating in the simplex mode in which the CPUs 12A... connect directly to the CPUs may drift with respect to each other.

Conversely, when the processing system 10 operates in a duplex mode (e.g., the CPUs operate in synchronized, lock-step operation), the clocks between routers 14 and the CPUs 12 to which they... ..not necessarily phase-locked).

The flow of data packets between the various elements of the processing system 10 is controlled by command symbols, which may appear at any time, even within... ..initiated by a CPU 12, or MP 18, and promulgated to all elements of the processing system 10 by the routers 14 to communicate an event requiring software action by all...command symbol is used in conjunction with near frequency operation as an aid to maintaining synchronization between the two clock signals that (1) transfer each symbol to, and load it in each receiving clock synchronization FIFO, and (2) that retrieves symbols from the FIFO.

SLEEP: This command symbol is sent by any element of the processing system 10 to indicate that no additional packet (after the one currently being transmitted, if... ..received.

SOFT RESET (SRST): The SRST command symbol is used as a trigger during the processes ("synchronization" and "reintegration," described below) that are used to synchronize symbol transfers between the CPUs 12 and the routers 14A, 14B, and then to place... ..SYNC command symbol is sent by a router 14 to the CPU 12 of the processing system 10 (i.e., the sub-processor systems 10A/10B) to establish frequency-lock synchronization between CPUs 12 and routers 14A, 14B prior to entering duplex mode, or when in duplex mode to request synchronization, as will be discussed more fully below. The SYNC command symbol is used in conjunction... ..or duplex to simplex), among other things, as discussed further below in the section on Synchronization and Reintegration.

THIS LINK BAD (TLB): When any system element receiving a symbol from a TNet link L (e.g., a router, a CPU, or an I/O unit) notes an error when receiving a command symbol or packet, it will send a TLB... ..identical pairs of symbols that are compared to one another when pulled from the clock synchronization FIFOs..The DVRG command symbol signals the CPU 12 that a mis-compare has been noted. When received by the CPUs, a divergence detection process is entered whereby a determination is made by the CPUs which CPU may be failing... ..command symbols described above operate to control message flow between the various elements of the processing system 10 (e.g., CPUs 12, router 14, and the like), using principally the BUSY however, an "end node" (i.e., a CPU 12 or I/O unit 17 - Fig. 1) may not assert backpressure because one of its transmit ports is backpressured... ..Improperly addressed packets are discarded by the router 14.

When a system element of the processing system 10 receives a BUSY command symbol on a TNet link L on which it... ..other command symbols (READY, BUSY, etc.).

Whenever a TNet port of an element of the processing system 10 detects receipt of a READY command symbol, it will terminate transmission of FILL... ..receives.

As will be seen, all elements (e.g., router 14, CPUs 12) of the processing system 10 that connect to a TNet link L for receiving transmitted symbols will receive those symbols via a clock synchronization (CS) FIFO. For example, as discussed above, the interface units 24 of CPUs 12 include all CS FIFOs 102x, 102y (illustrated in Fig. 6). The... ..depth to allow for speed matching, and the elastic FIFOs must provide sufficient depth for processing delays that may occur between transmission of a BUSY command symbol during receipt of a...another data byte in packet B. As packet A progresses to the next router, the process would be repeated. If the router 14 displaces more data bytes than the FIFO can... ..irrespective of its own findings.

SLEEP Protocol:

The SLEEP protocol is initiated by a maintenance processor via a maintenance interface (an on-line access port - OLAP), described below. The SLEEP protocol... ..reintegrate a slice of the system 10. Routers 14 must be idle (no packets in process) in order to change modes without causing data loss or corruption. When a SLEEP command symbol is received, the receiving element of processing system 10 inhibits initiation of transmission of any new packet on the associated transmit port... ..The HALT command symbol provides a mechanism for quickly informing

all CPUs 12 in a processing system 10 that is necessary to terminate I/O activity (i.e., message transmissions between... CPUs that receive HALT command symbols on either of their receive ports (of the interface units 24) will post an interrupt to the interrupt register 280 if the system halt interrupt... interrupt; Fig. 14A).

The CPUs 12 may be provided with the ability to disable HALT processing. Thus, for example, the configuration registers 75 of the interface units 24 can include a "halt enable register" that, when set to a predetermined state (eg., ZERO) disables HALT processing, but reporting detection of a HALT symbol as an error.

Router Architecture:

Referring now to... simplified block diagram of the router 14A is illustrated. The other routers 14 of the processing system 10 (e.g., routers 14B, 14', etc.) are of substantially identical construction and, therefore... these ports 4, 5 are structured to operate in a frequency locked environment when a processing system 10 is set for duplex mode operation. In addition, when in duplex mode, a... 5)) will receive the command/data symbols from the CPUs, pass them through the clock synchronization FIFOs 518 (discussed further below), and compare each symbol exiting the clock synchronization FIFOs with a gated compare circuit 517. When duplex operation is entered, a configuration register... 517 to activate the symbol by symbol comparison of the symbols emanating from the two synchronization FIFOs 518 of the router input logic 502 for the ports 4 and 5. Of... to that received, at substantially the same time, by the other port input.

To maintain synchronization in the duplex mode, the two port outputs of the router 14A that transmit to... mode, are duplicated by the routers 14, and returned to both CPUs.) The output logic units 504(sub(4)), 504(sub(5)) that are coupled directly to the CPUs 12 will... message packet identifies only one of the duplexed CPUs 12, e.g., CPU 12A) in synchronized fashion, presenting those symbols in substantially simultaneous fashion to the two CPUs 12. Of course, the CPUs 12 (more accurately, the associated interface units 24) receive the transmitted symbols with synchronizing FIFOs of substantially the same structure as that illustrated in Fig. 7A so that, even... from the FIFO structures by both CPUs 12 on the same instruction cycle, maintaining the synchronized, lock-step operation of the CPUs 12 required by the duplex operating mode.

As will... conjunction with configuration data written to registers contained in control logic 509 by the maintenance processor 18 (via the on-line access port 285' and serial bus 19A; see Fig. 1A... links L. The input logic 505 of each port input 502 also assists in maintaining synchronization - at least for those ports sending symbols in the near-frequency environment - by removing received... slower-receiving element receiving symbols from a faster-sending element could overload the input clock synchronization FIFO of the slower-receiving element. That is, if a slower clock is used to pull symbols from the clock synchronization FIFO put there by a faster clock, ultimately the clock synchronization FIFO will overflow.

The preferred technique employed here is to periodically insert SKIP symbols in... stream to avoid, or at least minimize, the possibility of an overflow of the clock synchronization FIFO (i.e., clock synchronization FIFO 518; Fig. 20A) of a router 14 (or CPU 12) due to a T... being slightly higher in frequency than the local clock used to pull symbols from the synchronization FIFO. Using SKIP symbols to by-pass a push (onto the FIFO) operation has the... stall each time a SKIP command symbol is received so that, insofar as the clock synchronization FIFO is concerned, the transmitting clock that accompanied the SKIP symbol was missing.

Thus, logic... the port inputs 502 will recognize, and key off receipt of, SKIP command symbols for synchronization in the near frequency clocking environment so that nothing is pushed onto the FIFO, but... 14, or between routers 14, or between a router 14 and an I/O interface unit 16A - Fig. 1) at a 50 Mhz rate, this allows for a worst case frequency... symbol by supplying FILL or IDLE symbols (which are received and pushed onto the clock synchronization FIFOs, but are not passed to the elastic FIFOs). In short, each elastic FIFO 506... received symbols are then communicated from the input register 516 and applied to a clock synchronization FIFO 518, also by the T(underscore)Clk. The clock synchronization FIFO 518 is logically the same as that illustrated in Figs. 8A and 8B, used in the interface units 24 of the CPUs 12. Here, as Fig. 20A shows, the clock synchronization FIFO

518 comprises a plurality of registers 520 that receive, in parallel, the output of... 516. Associated with each of the registers 520 is a two-stage validity (V) bit synchronizer 522, shown in greater detail in Fig. 20B, and discussed below. The content of each... 520, together with the one-bit content of each associated two-stage validity bit 5 synchronizer 522, are applied to a multiplexer 524, and the selected register/synchronizer pulled from the FIFO, and coupled to the elastic FIFO 506 by a pair of... is determined the state of the Push Select signal provided by a push pointer logic

unit 530; and, selection of which register 520 will supply its content, via the MUX 524... and loading of the register 520 selected by the push pointer logic 530. Similarly, the synchronization FIFO control logic 534 receives the clock signal local to the router (Rcv Clk) to... pointer logic 532.

Digressing for a moment, and referring to Fig. 20B, the validity bit synchronizer 522 is shown in greater detail as including a D-type flip-flop 541 with... 530 (Fig. 20A) selects the register 520 of the FIFO with which the validity bit synchronizer is associated for receipt of the next symbol - if not a SKIP symbol.

The delay... Truth Table, below). The D-type flip-flop 543 acts as an additional stage of synchronization, ensuring a stable level at the V output relative to the local Rec Clk. The... flip-flop 542, allowing the Pull signal (a periodic pulse from the sync FIFO Control unit 534) to clear the validity bit on this validity synchronizer 522 when the associated register 520 has been read. (Table omitted)

In summary, the validity synchronizer 522 operates to assert a "valid" (V) signal when a symbol is loaded in a... blocked from being routed out a particular port because another message is already in the process of being routed out that port. However, that other message in turn is also blocked... an incoming message packet bound for the CPUs will be replicated by the crossbar logic unit by routing the message packet to both port output 504(sub(4)) and 504(sub... P) identifies which of path (X or Y) should be used for accessing two sub-processing the device.

The routers 14 provide a capability of constructing a large, versatile routing network for, for example, massively parallel processing architectures. Routers are configured according to their location (i.e., level) in the network by...)) and 509(sub(k)) are such that bits "def" are used in the algorithmic process, then bits "abc" of the Region ID are compared to the content of the Device... the route to default register 509(sub(f)) to the final stage of the selection process: check logic 602. Check logic 602 operates to check the status of the port output... a lower level router, and may be located in one or another of the sub-processing systems 10A, 10B. Whether a router is an upper level or lower level router depends... of CPUs 12 and I/O devices 16 to one another, forming a massively parallel processing (MPP) system. Other such MPP systems may exist, and it is those routers configured as... captured. As soon as the message packet's Destination ID is so captured, the selection process begins, proceeding to the development of a target port address that will be used to... an error that will be posted to the MP18 via the router's (or interface unit's) OLAP for action.

Digressing, it should be appreciated that these protocol rules observed by the routers 14 are also observed by the CPUs 12 (i.e., interface units 24) and I/O packet interfaces 17.

Finally, when the router 14A is in the... directly with the CPUs 12A, 12B, and duplex mode is used, a duplex operation logic unit 638 is utilized to coordinate the port output connected to one of the CPUs 12A... was able to write instructions to the OLAP 285 that would be executed by the processors 20 to build a small memory image and routine to permit the CPU 12 to... the clock generation circuit design. There will be one clock generator circuit in each sub-processor system 10A/10B (Fig. 1) to maintain synchronism. Designated generally with the reference numeral 650... used by the various elements (e.g. CPU. 12, routers 14, etc.) of the sub-processor system containing the clock circuit 650 (e.g., 10A).

The clock generator 654 is shown... The 50 Mhz clock signals produced by the counter 663 are distributed throughout the sub-processor system where needed.

Turning now to Fig. 25, there is illustrated the interconnection and use... the clock circuits 650 used to develop synchronous clock signals for a pair of sub-processor systems 10A, 10B (Fig. 1) for frequency locked operation. As illustrated in Fig. 25, the two CPUs 12A and 12B of the sub-processor systems 10A, 10B each have a clock circuit 650, shown in Fig. 25 as clock... 654B of both CPUs 12. A driver and signal line 667 interconnects the two sub-processor systems to deliver the M(underscore)CLK signal developed by the oscillator circuit 652A to the clock generator 654B of the sub-processor system 10B. For fault isolation, and to maintain signal quality, the M(underscore)CLK signal is delivered to the clock generator 654A of the sub-processor system 10A through a separate driver and a loopback connection 668. The reason for the... the cable (not shown) will establish the connection shown in Fig. 25 between the sub-processor systems 10A, 10B; connected another way, the connections will be similar, but the oscillator 652B... Fig. 25, the M(underscore)CLK signal produced by the oscillator circuit 652A of sub-processing system 10A is used by both sub-processing systems 10A, 10B as their respective SYNC CLK signals and the various other clock signals... produced by the clock generators 654A, 654B. Thereby, the clock signals of the paired sub-processing systems 10A, 10B are synchronized for the frequency locked operation necessary for duplex mode.

The VCXOs 662 of the clock... This allows both clock generators 654A, 654B to continue to provide to the two sub-processing systems 10A, 10B clock signals in the face of improper operation of the oscillator circuit 652A, although the sub-processor systems may no longer be frequency-locked.

The LOCK signals asserted by the phase comparators... LOCK signal signifies that the 50 Mhz signals produced by a clock generator 654 are synchronized, both in phase and in frequency, to the M(underscore)CLK signal. Thus, if either... signal that accompanies the symbol stream, and is used to push symbols onto the clock synchronizing FIFO of the receiving element (router 14, or CPU 12) is substantially identical in frequency... not phase, to that of the receiving element used to pull symbols from the clock synchronization FIFOs. For example, referring to Fig. 23, which illustrates symbols being sent from the router... clock (Local Clk). The former (Rcv Clk) is used to push symbols onto the clock synchronization FIFOs 126 of each CPU, whereas the latter is used to pull symbols from the... much higher frequency clock signal. In such situations provision must be made to ensure that synchronization is maintained between the two CPUs as to symbols pulled from the clock synchronization FIFOs 126 of each.

Here, a constant ratio clocking mechanism is used to control operation of the two clock synchronization FIFOs 126, providing the clock signal that pulls symbols from the two FIFOs at the... control mechanism is shown, designated with the reference numeral 70. As Fig. 26A illustrates, clock synchronization FIFO control mechanism 700 includes an pre-settable, multi-stage serial shift register 702, the... ratio of the clock signal at which symbols are communicated and pushed onto the clock synchronization FIFOs 126 to the frequency of the clock signal used locally. Here, a 15 stages... that will be used as the Local Clk signal to pull symbols from the clock synchronization FIFOs 126, and to operate (update) the pull pointer counter 130. The selected output is... of the CPU 12 to the clock signal used to push symbols onto the clock synchronization FIFO 126, Rcv Clk, the serial shift register is preset so that M stages of... duplexed CPUs 12 with a 50 Mhz clock. Thus, symbols are pushed onto the clock synchronization FIFOs 126 of the CPUs at a 50 Mhz rate. Assume further that the clock... of the MUX 704, which produces the clock signal that pulls symbols from the clock synchronization FIFOs 126, Rcv Clk, will contain, for each 100 ns period, five clock pulses. Thus... five symbols will be pushed onto, and five symbols will be pulled from, the clock synchronization FIFOs 126.

This example is symbolically shown in Fig. 26B, while the timing diagram shown... labelled "IN" in Fig. 27) of the Rcv Clk will push symbols onto the clock synchronization FIFOs 126. During that same 100 ns period, the serial shift register 702 circulates a... clocks which would require additional storage (i.e., an increase in the size of the synchronization FIFO) and impose more latency.

The constant ratio clock circuit presented here (Figs. 26) is... frequency to a clock regime of a different, higher frequency. The use of a clock synchronization FIFO is necessary here for compensating effects of signal delays

when operating in synchronized, duplexed mode to receive pairs of identical command/data symbols from two different sources. However...so long as there are at least two registers in the place of the clock synchronization FIFO. Transferring data from a higher-frequency clock regime to a lower frequency clock regime... ...a wide range of possible clock ratios.

I/O Packet Interface:

Each of the sub-processor systems 10A, 10B, etc. will have some input/output capability, implemented with various peripheral units, although it is conceivable that the I/O of other sub-processor systems would be available so that a sub-processing system may not necessarily have local I/O. In any event, if local I/O... ...device (e.g., a signal line) would be received by the I/O packet interface unit 16 and used to form an interrupt packet that is sent to the CPU 12... ...OLAP bus, configuration information.

On-Line Access Port:

The MP 18 connects to the interface unit 24, memory controller (MC) 26, routers 14, and I/O packet interfaces with interface signals... ...OLAP 258 is essentially the same, regardless of what element (e.g. router 14, interface unit 24, etc.) it is used with. Fig. 28 diagrammatically illustrates the general structure of the... ...circuit chip used to implement certain of the elements discussed herein. For example, each interface

unit 24, memory controller 26, and router 14 is implemented by an application specific integrated circuit... ...of the OLAP 158 shown in Fig. 28 describes the OLAP associated with the interface unit 24, the MC 26, and the router 14 of the system.

As Fig. 28 shows...asymmetric variables, a "soft-vote" (SV) logic element 900 (Fig. 30A) is provided each interface unit 24 of each CPU 12. As Fig. 30 illustrates, the SV logic elements 900 of each interface unit 24 are connected to one another by a 2-bit SV bus 902, comprising bus lines 902a and 902b. Bus lines 902a carry one-bit values from the interface units 24 of CPU 12A to those of CPU 12B. Conversely, bus line 902b carries one... ...the CPU 12A.

Illustrated in Fig. 30B, is the SV logic element 900a of interface unit 24a of CPU 12A. Each SV logic element 900 is substantially identical in construction and... ...900a should be understood as applying equally to the other logic elements 900a (of interface unit 24b, CPU 12A), and 900b (of the interface units 24a, 24b of CPU 12B) unless noted otherwise. As Fig. 30B illustrates, the SV logic interface units 24a, 24b of the CPU 12A can communicate asymmetrical variables to each other.

In a... ...to the remote register 907 of logic element 902a (and that of the other interface unit 24b).

The logic elements 902 form a part of the configuration registers 74 (Fig. 5). Thus, they may be written by the processor unit(s) 20 by communicating the necessary data/address information over at least a portion of... ...local and remote registers 906 and 907.

The MUX 914 operates to provide each interface unit 24 of CPU 12A with selective use of the bus line 902a for the SV logic elements 900a, or for communicating a BUS ERROR signal if encountered during the reintegration process (described below) used to bring a pair of CPUs 12 into lock-step, duplex operation... ...same time, write the enable registers 912 of the logic element 900 of both interface units 24 of each CPU. One of the two logic elements 900 of each CPU will... ...it is the output enable registers 912 associated with the logic elements 900 of interface units 24a of both CPUs 12A, 12B that are written to enable the associated drivers 916. Thus, the output registers 904 of the interface units 24a of each CPU will be communicated to the bus lines 902; that is, the... ...to the bus line 902a, while the output register associated with logic element 900b, interface unit 24a of CPU 12B is communicated to bus line 902b. The CPUs 12 will both... ...again written by each CPU, followed again by reading the remote input registers 907. This process is repeated, one bit at a time, until the entire variable is communicated from the... ...each

CPU 12 to the remote input register of the other. Note that both interface units 24 of CPU 12B will receive the bit of asymmetric information.

One example of use... elements 900 are also used to communicate bus errors that may occur during the reintegration process to be described. When reintegration is being conducted, a REINT signal will be asserted. As... ERROR signal is selected by the MUX 914 and communicated to the bus line 902a.

Synchronization:

Proper operation of the sub-processing systems 10A, 10B (Figs. 1A, 2) whether operating independently (simplex mode), or paired and operating in synchronized lock-step (duplex mode), requires assurance that data communicated between the CPUs 12A, 12B and... routers 14A, 14B will be received properly, and that any initial content of the clock synchronization FIFOs 102 (of CPUs 12A, 12B; Fig. 5) and 519 (of routers 14A, 14B; Fig... erroneously interpreted as data or commands. The push and pull pointers of the various clock synchronization FIFOs 102 (in the CPUs 12) and 518 (in the routers 14) need to be... apart, and presetting the associated FIFO queues to some known state. This done, all clock synchronization FIFOs are initialized for near... in order to properly implement the lock-step operation of duplex mode operation, the clock synchronization FIFOs must be synchronized to operate with the particular source from which they receive data in order accommodate any... 14A, 14B to the CPUs 12A, 12B must be accounted for. It is the clock synchronization FIFOs 102 of the paired CPUs 12 that operate to receive message packet symbols, adjust... and present symbols to the two CPUs in a simultaneous manner to maintain lock-step synchronization necessary for duplex mode operation.

In similar fashion, each symbol received by the routers 14A... the CPUs (which is discussed further hereinafter). Again, it is the function of the clock synchronization FIFOs 518 of the routers 14A, 14B that receive message packets from the CPUs 12... so that the symbols received from the two CPUs 12 are retrieved from the clock synchronization FIFOs simultaneously.

Before discussing how the clock synchronization FIFOs of the CPUs and routers are reset, initialized, and synchronized, an understanding of their operation to maintain synchronous lock-step duplex mode operation is believed helpful. Thus, referring for the moment to Fig. 23, the clock synchronization FIFOs 102 of the CPUs 12A, 12B that receive data, for example, from the router... underscore)Clk, from the router 14A to the CPU 12B.

Consider operation of the clock synchronization FIFOs 102(sub(x)), 102(sub(y)), to receive identical symbol streams during duplex operation... held by the push and pull pointer counters 128, 130 for the CPU 12A (interface unit 24a), and the content of each of the four storage locations (byte 0. ..., byte 3... 6 show the same thing for the FIFO 102(sub(y)) of CPU 12B interface unit 24a for each symbol of the duplicated symbol stream.

Assuming the delay 640 is no...0" locations of the queues 126. This is because (1) the FIFOs 102 have been synchronized to operate in synchronism (a process described below), and (2) the push pointer counters 128 are clocked by the clock signal... of the symbol stream transmitted by the router 14A will be pulled from the clock synchronization FIFOs 102 of the CPUs 12A, 12B simultaneously, maintaining the required synchronization of received data when operating in duplex mode. In effect, the depths of the queues... order to achieve the operation just described with reference to Table 6, the reset and synchronization process shown in Fig 31A is used. The process not only initializes the clock synchronization FIFOs 102 of the CPUs 12A, 12B for duplex mode operation, but also operates to adjust the clock synchronization FIFOs 518 (Fig. 19A) of the CPU ports of each of the routers 14A, 14B for duplex operation. The reset and synchronization process uses the SYNC command symbol to initiate a time period, delineated by the SYNC CLK signal 970 (Fig. 31B), to reset and initialize the respective clock synchronization FIFOs of the CPUs 12A and 12B and routers 14A, 14B. (The SYNC CLK signal... It is of a lower frequency than that used to receive symbols by the clock synchronization FIFOs, T(underscore)Clk. For example, where T(underscore)Clk is approximately 50 MHz, the... signal is approximately 3.125 MHz.)

Turning now to Fig. 31A, the reset and initialization process begins at step 950 by switching the clock signals used by the CPUs 12A, 12B and routers 14A, 14B as the transmit (T(underscore)Clk) and the unit's local clock (Local Clk) clock signals so that they are derived from the same... In addition, configuration registers in the CPUs 12A, 12B (configuration registers 74 in the interface units 24) and the routers 14A, 14B (contained in control logic unit 509 of routers 14A, 14B) are set to the FreqLock state.

The following discussion involves step 952, and makes reference to the interface unit 24 (Fig.5), router 14A (Fig. 19A) and Figs. 31A and 31B. With the clock... otherwise be sent followed by a self-addressed message packet. Any message packet in the process of being received and retransmitted when the SLEEP command symbols are received and recognized by... per the destination address). The SLEEP command symbol operates to "quiece" router 14A for the synchronization process. The self-addressed message packet sent by the CPU 12A, when received back by the... message packet sent after the SLEEP command symbol would necessarily have to be the last processed by the router 14A.

At step 954 the CPU 12A checks to see if it...the router will assert a RESET signal 972 that is applied to the two clock synchronization FIFOs 518 contained in the input logic 505(sub(4)), 505(sub(5)) of the... receive symbols directly from CPUs 12A, 12B. RESET, while asserted, will hold the two clock synchronization FIFOs 518 in a temporarily non-operating reset state with the push and pull pointer... As each of the CPUs 12 receive SYNC symbols are detected by the storage and processing units of the packet receivers 96 (Figs. 5 and 6) cause the RESET signal to be asserted by the packet receivers 96 (actually, storage and processing elements 110; Fig. 6) of each CPU 12. the RESET signal is applied to the... t4), CPUs 12 and routers 14A, 14B de-assert the RESET signals, and the clock synchronization FIFOs of the CPUs 12A, 12, and routers 14A, 14B are released from their reset... the delay, the router 14A and CPUs 12 resume pulling data from their respective clock synchronization FIFOs and resume normal operation. The clock synchronization FIFOs of the router 14A begin pulling symbols from the queue (previously set by RESET... from the CPU 12A with the T(underscore)Clk will be pushed onto the clock synchronization FIFO at, for example, queue location 0 (or whatever other location pointed to by the... 0 (or whatever other location the push pointer was set to by RESET). The clock synchronization FIFOs of the router 14A are now synchronized to accommodate whatever delay 640 may be present in one communications path, relative to the... and the CPUs 12A, 12B.

Similarly, at the same virtual time, operation of the clock synchronization FIFOs 102 of both CPUs 12A, 12B is resumed, synchronizing them to the router 14A. Also, the CPUs 12A, 12B quit sending the SLEEP command... in favor of READY symbols, and resume message packet transmission, as appropriate.

That completes the synchronization process for the router 14A. However, the process must also be performed for the router 14B. Thus, the CPU 12A returns to step... however, assuming that the CPUs 12A, 12B are operating in duplex mode, the method and apparatus used to detect and handle a possible error, resulting in divergence of the CPUs from... via a message packet destined for a peripheral device of one or the other sub-processor systems 10A, 10B. Depending upon the destination of the outgoing message packet, step 1002 will... router 14 will issue an ERROR signal to the router control logic 509, causing the process to move to step 1004 where the router 14 detecting divergence will transmit a DVRG... time outs to occur. A router detecting divergence (without also detecting any simple link error) buys itself time to check the CRC of the received message packet by waiting for the... router 14, or received, all further message packets received from the CPUs and in the process of being routed when divergence was detected, or the DVRG symbol received, will be passed... 1010) contained in a one of the configuration registers 74 (Fig. 5) of the interface unit 24 of each CPU.

Returning for the moment to step 1006, the determination of which... local" is meant to refer to the router 14A, 14B contained in the same sub-processor system 10A, 10B as the CPU. For example, referring to Fig. 1A, router 14A is... bit mentioned above: the bit contained in one of the configuration registers 74 of interface unit 24(Fig. 5)

of each CPU. When set to a first state, that particular CPU...the other CPU. In response, the state machines (not shown) within the control and status unit 509 (Fig. 19A) changes the "favorite" bits described above.

A few examples may facilitate understanding... ..DVRG symbol will echo that symbol to the routers 14A, 14B, start its internal divergence process timer, and begin determination of whether to continue or terminate. Having received a TLB symbol...to diverge with no errors reported. This can happen only if software (running on the processors 20) uses known divergent data to alter state. For example, suppose each CPU 12 has... ..number of the CPU 12A will differ from that of the CPU 12B. If the processors use the serial number to change the sequence of instructions executed (say, by branching if the serial number comes after some value) or to modify the value contained in a processor register, the complete "state" of the CPUs 12 will differ. In such cases, the "asymmetrical... ..of the primary CPU simply allows one CPU, and thereby the system 10, to continue processing without software intervention.

- An error at the output of the interface unit 24 of a CPU 12 will be detected by the router 14A, 14B, depending upon... ..router 14A, 14B that connects to a CPU 12 will be detected by the interface unit 24 of the affected CPU. The CPU will send a TLB symbol to the faulty... ..possible failure and, without external intervention, and transparently to the system user, remove the failing unit (CPU 12A or 12B, or router 14A or 14B) from the system to obviate or... ..reintegration." The discussion will refer to the CPUs 12A, 12B, routers 14A, 14B, and maintenance processor 18A, 18B shown forming parts of the processing system 10 illustrated in Fig. 1A. In addition, discussion will refer to the processors 20a, 20b, the interface units 24a, 24b, and the memory controllers 26a, 26b (Fig. 2) of the CPUs 12A, 12B as single units, since that is the way they function.

Reintegration is used to place two CPUs in...both of the paired CPUs at virtually the same time.

The major steps in the process for changing from simplex mode operation of the one on-line CPU to duplex mode... ..greater detail by the flow diagrams of Figs. 33A - 33D, generally are:

1. Setup and synchronize the two CPUs (one on-line, the other off-line) and their connected routers to... ..the memory of the on-line CPU to the off-line CPU, maintaining a tracking process that monitors changes in the memory of the on-line CPU that have not been... ..and may need to be copied over to, the off-line CPU;
3. Setup and synchronize the CPUs to run a delayed (slave) duplex mode from the same instruction stream (lock... ..will write the predetermined registers (not shown) of the control registers 74 in the interface units 24 of CPUs 12A and 12B, to a next state (after a soft operation) in... ..the off-line CPU 12B.

Next, a sequence is entered (steps 1060 - 1070) that will synchronize the clock synchronization FIFOs of the CPUs 12A, 12B and routers 14A, 14B in much the same fashion... ..the same steps described above in connection with the discussion of Figs. 31A, 31B to synchronize the clock synchronization FIFOs. The on-line CPU 12A will send the sequence of a SLEEP symbol, self-addressed message packet, and SYNC symbol which, with the SYNC CLK signal, operates to synchronize CPUs and routers. Once so synchronized, the on-line CPU 12A then, at step 1066, sends a Soft Reset (SRST) command... ..of all configuration registers and control registers (e.g., configuration registers 74 of the interface units 24) cache, and the like to memory 28 of the on-line CPU, copying ...time to have the system 10 off-line for reintegration. For that reason, the reintegration process is performed in a manner that allows the on-line CPU to continue executing user... ..not match that of the off-line CPU. The reason for this is that normal processing by the processor 20 of the on-line CPU can change memory content after it has been copied... ..when a memory location is written in the on-line CPU 12A during the reintegration process it is marked as "dirty;" second, all copying of memory to the off-line CPU... ..may, however, limit the ability to detect two-bit errors. But, since the memory copying process will last for a only relatively short period of time, this risk is believed acceptable... ..memory location in CPU 12A is made (either an incoming I/O write, or a processor write operation). The returning data (that was copied over to the off-line CPU) would... ..controller 26 (Fig. 2) of the on-line CPU to monitor memory locations in the process of being copied over to the off-line CPU 12B. The memory controller uses

a...within the block had been written by another operation (e.g., a write by the processor 20, an I/O write, etc.), that prior write operation will flag the location in... ..still must be copied over to the off-line CPU 12B.

Returning to the reintegration process, and now to Fig. 33B, the memory tracking (AtomicWrite mechanism and using ECC to mark... ..entails writing a reintegration register (not shown; one of the configuration registers 74 of interface unit 24 - Fig. 5) to cause a reintegration (REINT) signal to be asserted. The REINT signal... ..is left alone. Throughout the incremental copy operations, the normal actions of the on-line processor will mark some memory locations dirty.

Several passes of incremental copying will need to be... ..the number of successful WriteConditional operations at the end of each pass through memory, the processors 20 can determine the effect of a given pass compared to the previous pass. When the benefits drop off, the processors 20 will give up on the precopy operations. At this point the reintegration process is ready to place the two CPUs 12A, 12B into lock-step operation.

Thus, the... ..in Fig. 33C, where at step 1100, the on-line CPU 12A momentarily halts foreground processing, i.e., execution of a user application. The remaining state (e.g., configuration registers, cache, etc.) of the on-line processors 20 and its caches is then read and written to a buffer (series of memory... ..to the off-line CPU 12B, together with a "reset vector" that will direct the processor units 20 of both CPUs 12A, 12B to a reset instruction.

Next, step 1106 will quiesce... ..to ensure that the FIFOs of the routers are clear, that the FIFOs of the processor interfaces 24 are clear, and no further incoming I/O message packets are forthcoming. At... ..symbol will be received and acted upon by both CPUs 12A, 12B, to cause the processor units 20 of each CPU to jump to the location in memory 28 containing the reset... ..a subroutine that will restore the stored state of both CPUs 12A, 12B to the processor units 20, caches 22, registers, etc. The CPUs 12A, 12B will then begin executing the same... ..enabling of the ECC bit to mark dirty locations must now be disabled, since the processors are doing the same thing to the same memory. During this stage of the reintegration... ..encountered by CPU 12A.

Meanwhile, the bus error in the CPU 12A will cause the processor unit 20 to be forced into an error-handling routine to determine (1) the cause of... ..error was caused by an attempt to read a memory location marked dirty. Accordingly, the processor unit 20 will initiate (via the BTE 88 -- Fig. 5) the AtomicWrite mechanism to copy the...the SRST symbols are now received by the CPUs 12A, 12B, they will cause both processor units 20 of the CPUs to be reset to start from the same location with the... ..will periodically update, e.g., a database or audit file that is indicative of the processing of the primary CPU up to that point in time of the update. Should the... ..in error-checking redundancy to the CPU 12B, in the same manner that the individual processor units 20a, 20b of the CPU 12A provide fail-fast, fault tolerance for the CPU - when... ..cost system is applicable, as illustrated in Fig. 34. As shown in Fig. 34, a processing system 10' includes the CPU 12A and routers 14A, 14B structured as described above. The... ..and the CPUs are also the same.

Thus, the CPU 12B' comprises only a single processor unit 20' and associated support components, including the cache 22', interface unit (IU) 24', memory controller 26', and memory 28'. Thus, while the CPU 12A is structured in the manner shown in Fig. 2, with cache processor unit, interface unit, and memory control redundancies, approximately one-half of those components are needed to implement CPU... ..stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... ..i.e., parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of... ..inputs connected to receive the data output from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... ..for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the... ..attached, without the assistance of any other CPU 12 in the system. Many prior

parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O device becomes... ..rest of the system. Other prior systems provide access to I/O through pairs of processors so that should one of the processors fail, access ...if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... ..multiprocessing system access to every peripheral, as done here, operates to extend the "primary"/"backup" process taught in the above-identified U.S. Patent No. 4,228,496. There, a multiple CPU system may have a primary process may running on one CPU, while a backup process resides in the background on another of the CPUs. Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be... ..e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system... ..CPU can access any I/O without the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... ..four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... ..an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time... ..indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... ..When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... ..response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Specification: ...of a given pass compared to the previous pass. When the benefits drop off, the processors 20 will give up on the precopy operations. At this point the reintegration process is ready to place the two CPUs 12A, 12B into lock-step operation.

Thus, the... ..in Fig. 33C, where at step 1100, the on-line CPU 12A momentarily halts foreground processing, i.e., execution of a user application. The remaining state (e.g., configuration registers, cache, etc.) of the on-line processors 20 and its caches is then read and written to a buffer (series of memory... ..to the off-line CPU 12B, together with a "reset vector" that will direct the processor units 20 of both CPUs 12A, 12B to a reset instruction.

Next, step 1106 will quiesce... ..to ensure that the FIFOs of the routers are clear, that the FIFOs of the processor interfaces 24 are clear, and no further incoming I/O message packets are forthcoming. At... ..symbol will be received and acted upon by both CPUs 12A, 12B, to cause the processor units 20 of each CPU to jump to the location in memory 28 containing the reset... ..a subroutine that will restore the stored state of both CPUs 12A, 12B to the processor units 20, caches 22, registers, etc. The CPUs 12A, 12B will then begin executing the same... ..enabling of the ECC bit to mark dirty locations must now be disabled, since the processors are doing the same thing ...encountered by CPU 12A.

Meanwhile, the bus error in the CPU 12A will cause the processor unit 20 to be forced-into an error-handling routine to determine (1) the cause of... ..error was caused by an attempt to read a memory location marked dirty. Accordingly, the processor unit 20 will initiate (via the BTE 88 -- Fig. 5) the AtomicWrite mechanism to copy the... ..the SRST symbols are now received by the CPUs 12A, 12B, they will cause both processor units 20 of the CPUs to be reset to start from the same location with the... ..will periodically update, e.g., a database or audit file that is indicative of the processing of the primary CPU up to that point in time of the update. Should the...in error-checking redundancy to the CPU 12B, in the same manner that the individual processor units 20a, 20b of the CPU 12A provide fail-fast, fault tolerance for the CPU - when... ..cost system is applicable , as illustrated in Fig. 34. As shown in Fig. 34, a processing system 10' includes the CPU 12A and routers 14A, 14B structured as described above. The... ..and the CPUs are also the same.

Thus, the CPU 12B' comprises only a single processor unit 20' and associated support components, including the cache 22', interface unit (IU) 24', memory controller 26', and memory 28'. Thus, while the CPU 12A is structured in the manner shown in Fig. 2, with cache processor unit, interface unit, and memory control redundancies, approximately one-half of those components are needed to implement CPU... ..stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... ..i.e, parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of... ..inputs connected to receive the data output from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... ..for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the... ..attached, without the assistance of any other CPU 12 in the system. Many prior parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O device becomes... ..rest of the system. Other prior systems provide access to I/O through pairs of processors so that should one of the processors fail, access to the corresponding I/O is still available through the remaining I/O... ..if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... ..multiprocessing system access to every peripheral, as done here, operates to extend the "primary"/"backup" process taught in the above-identified U.S. Patent No. 4,228,496. There, a multiple CPU system may have a primary process may running on one CPU, while a backup process resides in the background on another of the CPUs. Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be... ..e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system... ..CPU can access any I/O without the necessity of using the services of another processor . Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be...four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... ..an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time...indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... ..When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... ..response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Claims: ...A2

1. A multiple processing system, comprising:

a plurality of central processing units;

a plurality of input/output devices; and

a network interconnecting the central processing units and the input/output devices so that any one of the central processing units has communicative access to any one of the input/output devices without requiring use of any other of the plurality of central processing units.

Claims: ...dans lequel chacune des unites centrales de traitement est concue pour executer un ou plusieurs processus, comportant :

des moyens destines a executer un processus principal sur une premiere des unites centrales de traitement ; et

des moyens destines a executer un processus de reserve sur une seconde des unites centrales de traitement, le processus de reserve etant sensiblement identique, par sa mise en oeuvre et sa fonction, au processus principal.

9. Systeme selon la revendication 8, dans lequel, periodiquement au cours de l'execution du processus principal, la premiere unite centrale de traitement communique des donnees de message de point de... ..de message de point de controle etant indicatives de l'etat de l'execution du processus principal.

10. Systeme selon la revendication 9, dans lequel, dans le cas d'une defaillance... ..premiere unite centrale de traitement, la seconde unite centrale de traitement initie l'execution du processus de reserve associe pour prendre en charge la tache en cours d'accomplissement par le processus principal, comprenant l'extraction des donnees de message de point de controle memorisees a partir du dispositif d'entree/sortie predetermine, le processus de reserve devenant un nouveau processus principal.

11. Systeme selon la revendication 10, dans lequel une troisieme des unites centrales de traitement a, associe a elle, un nouveau processus de reserve qui est sensiblement identique, par sa mise en oeuvre et sa fonction, au nouveau processus principal.

12. Systeme selon la revendication 11, dans lequel, periodiquement au cours de l'execution du nouveau processus principal, la seconde unite centrale de traitement communique de nouvelles donnees de message de point... ..message de point de controle etant indicatives de l'etat de l'execution du nouveau processus principal.

Patent Assignee:

- Compaq Computer Corporation... ;
;;

Country	Number	Kind	Date	
Type	Pub. Date	Kind	Text	
...Compaq Computer Corporation (687797) 20555 SH 249 Houston, Texas 77070-2698 US	19			
Examination...	19			
Available Text	Language	Update	Word Count	
Total Word Count (Document A)				
Total Word Count (Document B)				
Total Word Count (All Documents)				

Specification: ...and 08/485,055 filed concurrently herewith.

The present invention is directed generally to data processing systems, and more particularly to a multiple processing system and a reliable system area network that provides connectivity for interprocessor and input/output... ...and communications systems to general purpose high availability commercial systems. The evolution of fault tolerant computers has been well documented (see D. P. Siewiorek, R. S. Swarz, "The Theory and Practice... ...and the Jet Propulsion laboratory began to apply fault tolerance to the development of guidance computers for aerospace applications. The 1960's also saw the development of the first AT&T electronic switching systems.

The first commercial fault tolerant machines were introduced by Tandem Computers in the 1970's for use in on-line transaction processing applications (J. Bartlett, "A NonStop Kernal," in proc. Eighth Symposium on Operating System Principles, pp... ...systems were introduced in the 1980's (O. Serlin, "Fault- Tolerant Systems in Commercial Applications," Computer, pp. 19-30, August 1984). Current commercial fault tolerant systems include distributed memory multi-processors, shared-memory transaction based systems, "pair-and-spare" hardware fault tolerant systems (see R. Freiburghouse, "Making Processing Fail-safe," Mini-micro Systems, pp. 255-264, May 1982; U.S. Patent No. 4... ...system.), and triple-modular-redundant systems such as the "Integrity" computing system manufactured by Tandem Computers Incorporated of Cupertino, California, assignee of this application and the invention disclosed herein.

Most applications of commercial fault tolerant computers fall into the category of on-line transaction processing. Financial institutions require high availability for electronic funds transfer, control of automatic teller machines, and... ...telecommunications systems.

Vendors of fault tolerant machines attempt to achieve both increased system availability, continuous processing, and correctness of data even in the presence of faults. Depending upon the particular system architecture, application software ("processes") running on the system either continue to run despite failures, or the processes are automatically restarted from a recent checkpoint when a fault is encountered. Some fault tolerant systems are provided with sufficient component redundancy to be able reconfigure around failed components, but processes running in the failed modules are lost. Vendors of commercial fault tolerant systems have extended fault tolerance beyond the processors and disks. To make large improvements in reliability, all sources of failure must be addressed... ...power supplies, fans and inter-module connections.

The "NonStop," and "Integrity" architectures manufactured by Tandem Computers Incorporated, (both respectively illustrated broadly in U.S. Patent No. 4,228,496 and U.S. Patent No. 4,278,496, both assigned to the assignee of this application; NonStop and Integrity are registered trademarks of Tandem Computers Incorporated) represent two current approaches to commercial fault tolerant computing. The NonStop system, as generally described in the above-identified U.S. Patent No. 4,278,496, employs an architecture that uses multiple processor systems designed to continue operation despite the failure of any single hardware component. In normal operation, each processor system uses its major components independently and concurrently, rather than as "hot backups". The NonStop system architecture may consist of up to 16 processor systems interconnected by a bus for interprocessor communication. Each processor system has its own memory which contains a copy of a message-based operating system. Each processor system controls one or more input/output (I/O) busses. Dual-porting of I/O controllers and devices provides multiple paths to each device. External storage (to the processor system), such as disk storage, may be mirrored to maintain redundant permanent data storage.

This... hardware, while fault recovery is the responsibility of the software.

Also, in the Nonstop multi-processor architecture, application software ("process") may run on the system under the operating system as "process-pairs," including a primary process and a backup process. The primary process runs on one of the multiple processors while the backup process runs on a different processor. The backup process is usually dormant, but periodically updates its state in response to checkpoint messages from the primary process. The content of a checkpoint message can take the form of complete state update, or currently most application code runs under transaction processing software which provides recovery through a combination of checkpoints and transaction two-phase commit protocols.

Interprocessor message traffic in the Tandem Nonstop architecture includes each processor periodically broadcasting an "I'm Alive" message for receipt by all the processors of the system, including itself, informing the other processors that the broadcasting processor is still functioning. When a processor fails, that failure will be announced and identified by the absence of the failed processor's periodic "I'm Alive" message. In response, the operating system will direct the appropriate backup processes to begin primary execution from the last checkpoint. New backup processes may be started in another processor, or the process may be run with no backup until the hardware has been repaired. U.S. Patent No. 4,278,496, example of this technique.

Each I/O controller is managed by one of the two processors to which it is attached. Management of the controller is periodically switched between the processors. If the managing processor fails, ownership of the controller is automatically switched to the other processor. If the controller fails, access to the data is maintained through another controller.

In addition to providing hardware fault tolerance, the processor pairs of the above-described architecture provide some measure of software fault tolerance. When a processor fails due to a software error, the backup processor frequently is able to successfully continue processing without encountering the same error. The software environment in the backup processor typically has different queue lengths, table sizes, and process mixes. Since most of the software bugs escaping the software quality assurance tests involve infrequent data dependent boundary conditions, the backup processes often succeed.

In contrast to the above-described architecture, the Integrity system illustrates another approach... fault recovery is the logical choice since few modifications to the software are required. The processors and local memories are configured using triple-modular-redundancy (TMR). All processors run the same code stream, but clocking of each module is independent to provide tolerance... three streams is asynchronous, and may drift several clock periods apart. The streams are re-synchronized periodically and during access of global memory. Voters on the TMR Controller boards detect and mask failures in a processor module. Memory is partitioned between the local memory on the triplicated processor boards and the global memory on the duplicated TMRC boards. The duplicated portions of the... techniques to detect failures. Each global memory is dual ported and is interfaced to the processors as well

to the I/O Processors (IOPs). Standard VME peripheral controllers are interfaced to a pair of busses through a Bus... the BIMs to switch control of all controllers to the remaining IOP. Mirrored disk storage units may be attached to two different VME controllers.

In the Integrity system all hardware failures... reintegrated on-line.

The preceding examples illustrate present approaches to incorporating fault tolerance into data processing systems. Approaches involving software recovery require less redundant hardware, and offer the potential for some... have been developed on other systems.

Thus, the systems described above provide fault tolerant data processing either by hardware (e.g, fail-functional, employing redundancy) or by software techniques (fail-fast... hardware). However, none of the systems described are believed capable of providing fault tolerant data processing, using both hardware (fail-functional) and software (fail-fast) approaches, by a single data processing system.

Computing systems, such as those described above, are often used for electronic commerce: electronic data interchange (EDI) and global messaging. Today's demands upon such electronic commerce, however, is demanding more and more throughput capacity as the number of users increases and... networks such as local area networks (LANS), and the like.

A key requirement for a server architecture is the ability to move massive quantities of data. The server should have high bandwidth that is scalable, so that added throughput capacity can be added... response time, latency affects service levels and employee productivity.

The present invention provides a multiple-processor system that combines both of the two above-described approaches to fault tolerant architecture, hardware redundancy and software recovery techniques, in a single system.

Broadly, the present invention includes a processing system composed of multiple sub-processing systems. Each sub-processing system has, as the main processing element, a central processing unit (CPU) that in turn comprises a pair of processors operating in lock-step, synchronized fashion ...execute each instruction of an instruction stream at the same time. Each of the sub-processing systems further include an input/output (I/O) system area network system that provides redundant communication paths between various components of the larger processing

system, including a CPU and assorted peripheral devices (e.g., mass storage

units, printers, and the like) of a sub-processing system, as well as between the sub-processors that may make up the larger overall processing system. Communication between any component of the processing system (e.g., a CPU and another CPU, or a CPU and any peripheral device, regardless of which sub-processing system it may belong to) is implemented by forming and transmitting packetized messages that are... responsible for choosing the proper or available communication paths from a transmitting component of the processing system to a destination component based upon information contained in the message packet. Thus, the... peripherals, but permits it to also be used for interprocessor communications.

As indicated above, the processing system of the present invention is structured to provide fault-tolerant operation through both "fail... at a variety of points in the various data paths between the (lock-step operated) processor elements of the CPU and its associated memory. In particular, the processing system of the present invention conducts error-checking at an interface, and in a manner... little impact on performance. Prior art systems typically implement error-checking by running pairs of processors, and checking (comparing) the data and instruction flow between the processors and a cache memory. This technique of error-checking tended to add delay to the... error-checking precluded use of off-the-shelf parts that may be available (i.e., processor/cache memory combinations on a single semiconductor chip or module). The present invention performs error-checking of the processors at points that operate at slower rates, such as the main memory and I/O interfaces which operate at slower speeds than the

processor-cache interface. In addition, the error-checking is performed at locations that allow detection of errors that may occur in the processors, their cache memory, and the I/O and memory interfaces. This allows simpler designs for... other data integrity checks.

Error-checking of the communication flow between the components of the processing system is achieved by adding a cyclic-redundancy-check (CRC) to the message packets that... "Good" (TPG) or "This Packet Bad" (TPB) - is appended to every packet. A maintenance diagnostic processor can use this information to isolate a link or router element that introduces an error... of topologies, so that alternate paths can be provided between any two elements of a processing system (e.g., between a CPU and an I/O device), for communication in the... so (e.g., by creating a "deadlock" condition, discussed further below).

The CPUs of a processing system are capable of operating in one of two basic modes: a "simplex mode" in... independently of the other, or a "duplex" mode in which pairs of CPUs operate in synchronized, lock-step fashion. Simplex mode operation provides the capability of recovering from faults that are... U.S. Pat. No. 4,228,496 which teaches a multiprocessing system in which each processor has the capability of checking on the operability of its sibling processors, and of taking over the processing of a processor found or believed to have failed). When operating in duplex mode, the paired CPUs both... fault tolerant platform for less robust operating systems (e.g., the UNIX operating system). The processing system of the present invention, with the paired, lock-step CPUs, is structured so that... masked (i.e., operating despite the existence of a fault), primarily through hardware.

When the processing system is operating in duplex mode, each CPU pair uses the I/O system to access any peripheral of the processing system, regardless of which (of the two, or more) sub-processor system the peripheral may be ostensibly a member of. Also, in duplex mode, message packets... message for the CPU pair (from either a peripheral device such as a mass storage unit or from a processing unit), will replicate the message and deliver it to both CPUs of the pair using synchronization methods that ensure that the CPUs remain synchronized. In effect, the duplex CPU pair, as viewed from the I/O system and other... as a single CPU. Thus, the I/O system, which includes elements from all sub-processing systems, is made to be seen by the duplex CPU pair as one homogeneous system... a multiprocessor system in which the CPU of any one is actually a pair of synchronized, lock-step CPUs.

Yet another important aspect of the present invention is that interrupts issuing... interrupts via the message packet system ensures that they will arrive at duplexed CPUs in synchronized fashion, in the same manner as I/O message packets. Interrupt message packets will contain... the system. In addition, using the same messaging system to communicate data between I/O units and the CPUs and to communicate interrupts to the CPUs preserves the ordering of I... the implementation of a technique of validating access to the memory of any CPU. The processing system, as structured according to the present invention, permits the memory of any CPU to a CPU and any other component of the processor system. Thereby, the individual processor units of the CPU are removed from the more mundane tasks of getting information from memory and out onto the TNet network, or accepting information from the network. The processor unit of the CPU merely sets up data structures in memory containing the data to be... is required, where in memory the response is to be placed when received. When the processor unit completes the task of creating the data structure, the block transfer engine is notified to... response is received, it is routed to the expected memory location identified, and notifies the processor unit that the response was received.

Further aspects and features of the present invention will become... invention, which should be taken in conjunction with the accompanying drawings.

Fig. 1A illustrates a processing system constructed in accordance with the teachings of the present invention, and Figs. 1B and 1C illustrate two alternate configurations of the processing system of Fig. 1A, employing clusters or arrangements of the processing system of Fig. 1A;

Fig. 2 illustrates, in simplified block diagram form, the central processing unit (CPU) that forms a part of each sub-processor system of Figs. 1A - 1C;

Figs. 3A - 3D and 4A - 4C illustrate the construction of... ..the area network I/O system shown in Fig. 2;

Fig. 5 illustrates the interface unit that forms a part of the CPUs of Fig. 2 to interface the processor and memory with the I/O area network system;

Fig. 6 is a block diagram, illustrating a portion of packet receiver of the interface unit of Fig. 5;

Fig. 7A diagrammatically illustrates the clock synchronization FIFO (CS FIFO) used by the packet receiver section packet receiver shown in Fig. 6;

Fig. 7B is an block diagram of a construction of the clock synchronization FIFO structure shown in Fig. 7A;

Fig. 8 illustrates the cross-connections for error-checking outbound transmissions from the two interface units of a CPU;

Fig. 9 illustrates an encoded (8B to 9B) data/command symbol;

Fig. 10 illustrates the method and structure used by the interface unit of Fig. 5 to cross-check for errors data being transferred to the memory controllers... ..of a CPU of Fig. 2 to other (external to the CPU) components of the processing system;

Fig. 12 is a block diagram that diagrammatically illustrates the formation of an address... ..14A illustrates the logic for posting interrupt requests to queues in memory and to the processor units of the CPU of Fig. 2;

Fig. 14B illustrates the process used to form a memory address for a queue entry;

Fig. 15 is a block... ..data output constructs formed in the memory of the CPU of Fig. 2 by a processor unit, and containing data to be sent via the area I/O networks shown in Figs. 1A - 1C, and also illustrating the block transfer engine (BTE) unit of the interface unit of Fig. 5 that operates to access the data output constructs for transmission to the... ..pair of memory controllers between memory of a CPU of Fig. 2 and its interface unit for accessing from memory 72 bits of data, including two simultaneously-accessed 32-bit words... ..other for error-checking;

Fig. 19A is a simplified block diagram illustration of the router unit used in the area input/output networks of the processing systems shown in Figs. 1A - 1C;

Fig. 19B illustrates comparison on two port inputs of the router unit of Fig. 19A;

Fig. 20A is a block diagram the construction of one of the six input ports of the router unit shown in Fig. 19A;

Fig. 20B is a block diagram of the synchronization logic used to validate command/data symbols received at an input port of the router unit of Fig. 19A;

Fig. 21A is a block diagram illustration of the target port selection... ..is a block diagram illustration of one of the six output ports of the router unit shown in Fig. 19A;

Fig. 23 is an illustration of the method used to transmit identical information to a duplexed pair CPUs of Fig. 2 in synchronized fashion when the processing system is operating in lock-step (duplex) mode, using a pair the FIFOs of Fig... ..is a simplified block diagram illustrating the clock generation system of each of the sub-processing systems of Figs. 1A - 1C for developing the plurality of clock signals used to operate the various elements of that sub-processing system;

Fig. 25 illustrates the topology used to interconnect the clock generation systems of paired sub-processing systems for synchronizing the various clock signals of the pair of sub-processing systems to one another;

Fig. 26A and 26B illustrates a FIFO constant rate clock control logic used to control the clock synchronization FIFO of Figs. 8 or 20 in the situation when the two clocks used to... ..structure of the on-line access port (OLAP) used to provide access to the maintenance

processor (MP) to the various elements of the system of Fig. 1A (or those of Figs... ..the soft-flag logic used to handle asymmetric variables between the CPUs of paired sub-processor systems operating in duplex mode;

Fig. 31A shows a flow diagram, and Fig. 31B illustrates a portion of SYNC CLK, both of which are used to reset and synchronize the clock synchronization FIFOs of the CPUs and routers of the processing system of Fig. 1A that receive information from each other;

Fig. 32 is a flow... ..33A - 33D generally illustrate the procedure used to bring an one of the CPUs of processing system shown in Fig. 1A into lock-step, duplex mode operation with the other of the CPUs without measurably halting operation of the processing system; and

Fig. 34 illustrates a reduced cost architecture incorporating teachings of the invention; and... ..to the figures and, for the moment, principally Fig. 1A, there is illustrated a data processing system, designated with the reference 10, constructed according to the various teachings of the present invention. As Fig. 1A shows, the data processing system 10 comprises two sub-processor systems 10A and 10B each of which are substantially the same in structure and function... ..should be appreciated that, unless noted otherwise, a description of any one of the sub-processor systems 10 will apply equally to any other sub-processor system 10.

Continuing with Fig. 1A therefore, each of the sub-processor systems 10A, 10B is illustrated as including a central processing unit (CPU) 12, a router 14, and a plurality of input/output (I/O) packet interfaces... ..one of the I/O packet interfaces 16 will also have coupled thereto a maintenance processor (MP) 18.

The MP 18 of each sub-processor system 10A, 10B connects to each of the elements of that sub-processor system via an IEEE 1149.1 test bus 17 (shown in phantom in Fig. 1A... ..accompanying clock signal. As Fig. 1A further illustrates, TNet Links L also interconnect the sub-processor systems 10A and 10B to one another, providing each sub-processor system 10 with access to the I/O devices of the other as well as inter-CPU communication. As will be seen, any CPU 12 of the processing system 10 can be given access to the memory of any other CPU 12, although... ..the memory of a CPU 12 by a wayward peripheral device 17.

Preferably, the sub-processor systems 10A/10B are paired as illustrated in Fig. 1A (and Figs 1B and 1C, discussed below), and each sub-processor system 10A/10B pair (i.e., comprising a CPU 12, at least one router 14... ..12A) connects, by a TNet Link L to a router (14A) of the corresponding sub-processor system (e.g., 10A). Conversely, the Y port connects the CPU (12A) to the router (14B) of the companion sub- processor system (10B). This latter connection not only provides a communication path for access by a CPU (12A) to the I/O devices of the other sub-processor system (10B), but also to the CPU (12B) of that system for inter-CPU communication.

Information is communicated between any element of the processing system 10 and any other element (e.g., CPU 12A of sub-processor system 10A) of the system and any other element of the system (e.g., an I/O device associated with an I/O packet interface 16B of sub-processor system 10B) via message "packets." Each message packet is made up of a number of... ..this reason, a unique method of receiving the symbols at the receiver, using a clock synchronization first-in-first-out (CS FIFO) storage structure (described more fully below), has been developed... ..operation means just that: the frequencies of the clock signals of the transmitter and receiver units are locked, although not necessarily in phase. Frequency locked clock signals are used to transmit symbols between the routers 14A, 14B and the CPUs 12 of paired sub-processor systems (e.g., sub-processor systems 10A, 10B, Fig. 1A). Since the clocks of the transmitting and receiving element are not phase related, a clock synchronization FIFO is again used -- albeit operating in a slightly different mode from that used for... ..difference, as will be seen, is due to the fact that pairs of the sub-processor systems 10 can be operated in a synchronized, lock-step mode, called duplex mode, in which each CPU 12 operates to execute the... ..1A illustrates another feature of the invention: a cross-link connection between the two sub-processor systems 10A, 10B through the use of additional routers 14 (identified in Fig. 1A as... ..RY(sub(1)), and RY(sub(2)) form a cross-link connection between the sub-processors 10A, 10B (or,

as shown, "sides" X and Y, respectively) to couple them to I... ..the routers RX(sub(2)) and RY(sub(2)) provide the I/O packet interface units 16x and 16y with a dual ported interface. Of course, it will now be evident... ..lend themselves to being used in a manner that can extend the configuration of the processing system 10 to include additional sub-processor systems such as illustrated in Figs. 1B and 1C. In Fig. 1B, for example, one... ..of each of the routers 14A and 14B is used to connect the corresponding sub-processor systems 10A and 10B to additional sub-processor systems 10A' and 10B' forming thereby a larger processing system comprising clusters of the basic processing system 10 of Fig. 1.

Similarly, in Fig. 1C the above concept is extended to form an eight sub-processor system cluster, comprising sub-processor systems pairs 10A/10B, 10A'/10B', 10A''/10B'', and 10A'''/10B'''. In turn, each of the sub-processor systems (e.g., sub-processor system 10A) will have essentially the same basic minimum configuration of a CPU 12, a... ..by a I/O packet interface 16, except that, as Fig. 1C shows, the sub-processor systems 10A and 10B include additional routers 14C and 14D, respectively, in order to extend the cluster beyond sub-processor systems 10A'/10B' to the sub-processor systems 10A''/10B'' and 10A'''/10B'''. As Fig. 1C further illustrates, unused ports 4 and... ..the routers 14 when configuring the topology of the system 10, any CPU 12 of processing system 10 of Fig. 1C can access any other "end unit" (e.g., a CPU or I/O device) of any of the other sub-processor systems. Two paths are available from any CPU 12 to the last router 14 connecting to the I/O packet interface 16. For example, the CPU 12B of the sub-processor system 10B' can access the I/O 16''' of sub-processor system 10A''' via router 14B (of sub-processor system 10B'), router 14D, and router 14B (of sub-system 10B'') and, via link LA... ..10A'''), OR via router 14A (of sub-system 10A'), router 14C, and router 14A (sub-processor system 10A'''). Similarly, CPU 12A of sub-processor system 10A'' may access (via two paths) memory contained in the CPU 12B of sub-processor 10B to read or write data. (Memory accesses by one CPU 12 of another component of the processing system requires, as will be seen, the components seeking access to have authorization to do... ..prevents corruption of memory data of a CPU by erroneous access.)

The topology of the processing system shown in Fig. 1B is achieved by using port 1 of the routers 14A, 14B, and auxiliary TNet links LA, to connect to the routers 14A', 14B' of sub-processor systems 10A', 10B'. The topology thereby obtained establishes redundant communication paths between any CPU 12 (12A, 12B, 12A', 12B') and any I/O packet interface 16 of the processing system 10 shown in Fig. 1B. For example, the CPU 12A' of the sub-processor system 10A' may access the I/O 16A of sub-processor system 10A by a first path formed by the router 14A' (in port 4, out... ..shown in Fig. 1B. By interconnecting one port of each router 14 of each sub-processor pair, and using additional auxiliary TNet links LA (illustrated in Fig. 1C with the dotted line connections) between the ports 1 of the routers 14 (14A'' and 14B'') of sub-processor systems 10A'', 10B'' and 10A''', 10B''', two separate, independent data paths can be found between any CPU 12 and any I/O packet interface 16. In this fashion, any end unit (i.e., a CPU 12 or an I/O packet interface 16) will have at least two paths to any other end unit.

Providing alternate paths of access between any two end units (e.g., between a CPU 12 and any other CPU 12, or between any CPU... ..any two of the remaining fault domains. Here, a fault domain could be a sub-processor system (e.g., 10A). Thus, if the sub-processor system 10A were brought down because of a failure the electrical power being supplied, without... ..TNet link LA between the routers 14A''' and 14B''', the CPU 12B of the sub-processor system 10B would have lost access to the I/O packet interface 16''' (via router... ..with the loss of the router 14A (and router 14C) by loss of the sub-processor system 10A, communications between the CPU 12B is still possible via the route of router... ..equally to CPU 12B. As Fig. 2 shows, the CPU 12A includes a pair of processor units 20a, 20b that are configured for synchronized, lock-step operation in that both processor units 20a, 20b receive and execute identical instructions, and issue identical data and command outputs, at substantially the same moments in time. Each of the processor units 20a and 20b is connected, by a bus 21 (21a, 21b) to a corresponding cache memory 22. The particular type of processor units used could contain sufficient internal cache memory so that the cache memory 22 would not... ..22 could be used to supplement any cache memory that may be internal to the processor units 20. In any event, if the cache memory 22 is used, the bus 21 is... ..22 address bits, 3 bits of parity covering the address, and 7 control bits.

The processors 20a, 20b are also respectively coupled, via a separate 64-bit address/data bus 23 to X and Y interface units 24a, 24b. If desired, the address/data communicated on each bus 23a, 23b could also be protected by parity, although this will increase the width of the bus. (Preferably, the processors 20 are constructed to include RISC R4000 type microprocessors, such as are available from the MIPS Division of Silicon Graphics, Inc. of Santa Clara, California.)

The X and Y interface units 24a, 24b operate to communicate data and command signals between the processor units 20a, 20b and a memory system of the CPU 12A, comprising a memory controller (MC... ..MC halves 26a and 26b) and a dynamic random access memory array 28. The interface units 24 interconnect to each other and to the MCs 26a, 26b by a 72-bit... ..accompanied by 8 bits of ECC) are written to the memory 28 by the interface units 24, one interface unit 24 will drive only one word (e.g., the 32 most significant portion) of the doubleword being written while the other interface unit 24 writes the other word of the double word (e.g., the least significant 32-bit portion of the doubleword). In addition, on each write operation the interface units 24a, 24b perform a cross-check operation on the data not written by that interface unit 24 with the data written by the other to check for errors; on read operations... ..accessed corresponds to the address of the location from which the doubleword was stored.

Interface units 24a, 24b of the CPU 12A form the circuitry to respectively service the X and Y (I/O) ports of the CPU 12A. Thus, the X interface unit 24a connects by the bi-directional TNet Link Lx to a port of the router 14A of the processor system 10A (Fig. 1A) while the Y interface unit 24b similarly connects to the router 14B of the processor system 10B by TNet Link Ly. The X interface unit 24a handles all I/O traffic between the router 14A and the CPU 12A of the sub-processor system 10A. Likewise, the Y interface unit 24b is responsible for all I/O traffic between the CPU 12A and the router 14B of companion sub-processor system 10B.

The TNet Link Lx connecting the X interface unit 24a to the router 14A (Fig. 1) comprises, as above indicated, two 10-bit buses... ..sub(x)) carries data incoming from the router 14A. In similar fashion, the Y interface unit 24b is connected to the router 14B (of the sub-processor system 10B) by two 10-bit buses: 30(sub(y)) (for outgoing transmissions) and 32... ..y)) (for incoming transmissions), together forming the TNet Link Ly.

The X and Y interface units 24a, 24b are synchronously operated in lock-step, performing substantially the same operations at substantially the same times. Thus, although only the X interface unit 24a actually transmits data onto the bus 30(sub(x)), the same output data is being produced by the Y interface unit 24b, and used for error-checking. The Y interface unit 24b output data is coupled to the X interface unit 24a by a cross-link 34(sub(y)) where it is received by the X interface unit 24a and compared against the same output data produced by the X interface unit. In this way the outgoing data made available at the X port of the CPU... ..the port of the CPU 12A is checked. The output data from the Y interface unit 24b is coupled to the Y port by a 10-bit bus 30(sub(y)), and also to the X interface unit 24a by the 9-bit cross-link 34(sub(y)) where is checked with that produced by the X interface unit.

As mentioned, the two interface units 24a, 24b operate in synchronous, lock-step with one another, each performing substantially the same... ..X and/or Y ports of the CPU 12A must be received by both interface units 24a, 24b to maintain the two interface units in this lock-step mode. Thus, data received by one interface unit 24a, 24b is passed to the other, as indicated by the dotted lines and 9... ..sub(x)) (communicating incoming data being received at the X port by the X interface unit 24a to the Y interface unit 24b) and 36(sub(y)) (communicating data received at the Y port by the Y interface unit 24b to the X interface unit 24a).

Certain more robust operating systems are structured with a fault-tolerant capability in the... ..example, U.S. Patent No. 4,817,091 teaches a multiprocessor system in which each processor periodically messages each of the processors of the system (including itself), under software control, to thereby provide an indication of continuing operation. Each of the processors, in addition to performing its normal tasks, operates as a backup processor to another of the processors. In the event one of the backup processors fails to receive the messaged indication from a sibling processor, it will take over the operation of that sibling (now thought to be inoperative), in... ..platform for

both types of software. Thus, when a robust operating system is available, the processing system 10 can be configured to operate in a "simplex" mode in which each of... ..left, in most instances, to software.

Alternatively, for less robust operating systems and software, the processing system 10 provides a hardware-based fault-tolerance by being configured to operate in a... ..g., CPUs 12A, 12B) are coupled together as shown in Fig. 1A, to operate in synchronized, lock-step fashion, executing the same instructions at the substantially the same moment in time... ..data and command symbols. In order to simplify the design of the CPU 12, the processors 20 are precluded from communicating directly with any outside entity (e.g., another CPU 12... ..O device via the I/O packet interface 16). Rather, as will be seen, the processor will construct a data structure in memory and turn over control to the interface units 24. Each interface unit 24 includes a block transfer engine (BTE; Fig. 5) configured to provide a form of... ..to the destination according to information contained in the message packet.

The design of the processing system 10 permits a memory 28 of a CPU to be read or written by... ..via the routers 14. Accordingly, before continuing with the description of the construction of the processing system 10, it would be of advantage to understand first the configuration of the data...information.

As indicated, the HADC message packet operates to communicate write data between the end units (e.g., CPU 12) of the processing system 10. Other message packets, however, may be differently constructed because of their function and... ..CRC. The HC message packet is used to acknowledge a request to write data.

Interface Unit:

The X and Y interface units 24 (i.e., 24a and 24b - Fig. 2) operate to perform three major functions within the CPU 12: to interface the processors 20 to the memory 28; to provide an I/O service that operates transparently to, but under the control of, the processors; and to validate requests for access to the memory 28 from outside sources.

Regarding first the interface function, the X and Y interface units 24a, 24b operate to respectively communicate processors 20a, 20b to the memory controllers (Mcs 26a, 26b) and memory 28 for writing and... ..fast checking of the data read/written. For example, write operations have the two interface units 24a, 24b cooperating to cross-check the data to be written to ensure its integrity (and at the same time, the interface units 24 will operate) to develop an error correcting code (ECC) that covers, as will be appropriate address.

With respect to I/O access, the processors 20 are not provided with the ability to communicate directly with the input/output systems... ..must write data structures to the memory 28 and then pass control to the interface units 24 which perform a direct memory access (DMA) operation to retrieve those data structures, and... ..indicated in the data structure itself.)

The third function of the X and Y interface units 24, access validation to the memory 28, uses an address validation and translation (AVT) table maintained by the interface units. The AVT table contains an address for each system component (e.g., an I/O... ..the incoming message packets are virtual addresses. These virtual addresses are translated by the interface unit to physical addresses recognizable by the memory control units 26 for accessing the memory 28.

Referring to Fig. 5, illustrated is a simplified block diagram of the X interface unit 24a of the CPU 12A. The companion Y interface unit 24b (as well as the interface units 24 of the CPU 12B, or any other CPU 12) is of substantially identical construction. Accordingly, it will be understood that a description of the interface unit 24a will apply equally to the other interface units 24 of the processing system 10.

As Fig. 5 illustrates, the X interface unit 24a includes a processor interface 60, a memory interface 70, interrupt logic 86, a block transfer engine (BTE) 88, access validation and translation logic 90, a packet transmitter 94, and a packet receiver 96.

Processor Interface:

The processor interface 60 handles the information flow (data and commands) between the processor 20a and the X interface unit 24a. A processor bus 23, including a 64 bit address and data bus (SysAD) 23a and a 9 bit command bus 23b, couples the processor 20a and the processor interface 60 to one another. While the SysAD bus 23a carries memory address and data... ..and qualifying commands carried at substantially the same time on the SysAD bus 23a. The processor interface 60 operates to interpret commands issued by the processor unit 20a in order to pass reads/writes to memory or control registers of the processor interface. In addition, the processor interface 60 contains temporary storage (not shown) for buffering addresses and data for access to... ..26). Data and command information read from memory is similarly buffered en route to the processor unit 20a, and made available when the processor unit is ready to accept it. Further, the processor interface 60 will operate to generate the necessary interrupt signalling for the X interface unit 24a.

The processor interface 60 is connected to a memory interface 70 and to configuration registers 74 by a bi-directional 64 bit processor address/data bus 76. The configuration registers 74 are a symbolic representation of the various control registers contained in other components of the X interface unit 24a, and will be discussed when those particular components are discussed. However, although not specifically... ..throughout other of the logic that is used to implement the X interface 24a, the

processor address/data bus 76 is likewise coupled to read or write to those registers.

Configuration registers 74 are read/write accessible to the processor 20a; they allow the X interface unit to be "personalized." For example, one register identifies the node address of the CPU 12A... ..with the CPU 12A; another, readable only, contains a fixed identification number of the interface unit 24, and still other registers define areas of memory that can be used by, for... ..logic 90, etc.) employing them are discussed.

The memory interface 70 couples the X interface unit 24a to the memory controllers 26 (and to the Y interface unit 24b; see fig. 2) by a bus 25 that includes two 36 bi-directional bit... ..25a, 25b. The memory interface operates to arbitrate between requests for memory access from the processor unit 20, the BTE 88, and the AVT logic 90. In addition to memory accesses from the processor unit 20a, the memory 28 may also be accessed by components of the processing system 10 to, for example, store data requested to be read by the processor unit 20a from an I/O unit 17, or memory 28 may also be accessed for I/O data structures previously set up in memory by the processor unit. Since these accesses are all asynchronous, they must be arbitrated, and the memory interface 70... ..command information accessed from the memory 28 is coupled from the memory interface to the processor interface 60 by a memory read bus 82, as well as to an interrupt logic... ..doubleword quantities. However, while the memory interfaces 70 of both the X and Y interface units 24a ...by the memory interface 70 are coupled to the memory interface by the companion interface unit 24 where they are compared with the same 32 bits for error.

Digressing for the... ..containing interrupt information are received, that information is conveyed to the interrupt logic 86 for processing and posting for action by the processor 20, along with any interrupts generated internal to the CPU 12A. Internally generated interrupts will... ..register 71 (internal to the interrupt logic 86), indicating the cause of the interrupt. The processor 20 can then read and act upon the interrupt. The interrupt logic is discussed more fully below.

The BTE 88 of the X interface unit 24a operates to perform direct memory accesses, and provides the mechanism that allows the processors 20 to access external resources. The BTE 88 can be set-up by the processors 20 to generate I/O requests, transparent to the processors 20 and notify the processors when the requests are complete. The BTE logic 88 is discussed further below.

Requests for... ..8 byte wide format necessary for storing in the memory 28.

Outgoing message packets containing processor originated transaction requests (e.g., a read request asking for a block data from an I/O unit) are monitored by the request transaction logic (RTL) 100. The RTL 100 provides a

time... will generate an interrupt (handled and reported by the interrupt logic 86) to inform the processor 20 that the request was not honored. In addition, the RTL 100 will validate responses... 28 (by the DMA operation of the BTE 86) at a location known to the processor 20 so that it can locate the response.

Each of the CPUs 12 are checked... discussed. One such check is an on-going monitor of the operation of the interface units 24a, 24b of each CPU. Since the interface units 24a, 24b operate in lock-step synchronism checking can be performed by monitoring the operating states of the paired interface units 24a, 24b by a continuous comparison of certain of their internal states. This approach is implemented by using one stage of a state machine (not shown) contained in the unit 24a of CPU 12A, and comparing each state assumed by that stage with its identical state machine stage in the interface unit 24b. All units of the interface units 24 use state machines to control their operations. Preferably, therefore, a state machine of the memory interface 70 that controls the data transfers between the interface unit 24 and the MC 26 is used. Thus, a selected stage of the state machine used in the memory interface 70 of the interface unit 24a is selected. An identical stage of a state machine of one of the interface unit 24b is also selected. The two selected stages are communicated between the interface units 24a, 24b and received by a compare circuit contained in both interface units 24a, 24b. As the interface units operate lock-step with one another, the state machines will likewise march through the same identical states, assuming each state at substantially the same moments in time. If an interface unit encounters an error, or fails, that activity will cause the interface units to diverge, and the state machines will assume different states. The time will come when... that will bring to the attention of the CPUs 12A (or 12B) that the interface units 24a, 24b of that CPU are no longer in lock-step, and to act accordingly... X port, receiving only those message packets transmitted by the router 14A of the sub-processor system 10A (Fig. 1A). The Y port is serviced by the Y interface unit 24b to receive message packets from the router 14B of the companion sub-processor system 10B. However, both interfaces (as well as Mcs 26 and processor 20), as has been indicated, are basically mirror images of one another in that both... in both structure and function. For this reason, message packet information, received by one interface unit (e.g., 24a) must be passed for processing also to the companion interface unit (e.g., 24b). Further, since both interface units 24a, 24b will assemble the same message packets for transmission from the X or the Y ports, the message packet being transmitted by the interface unit (e.g., 24b) actually being communicated from the associated port (e.g., the Y port) will also be coupled to the other interface unit (e.g., 24a) for cross-checking for errors. These features are illustrated in Figs. 6... receiving portions of the packet receivers 96 (96x, 96y) of the X and Y interface units 24a, 24b are broadly illustrated. As shown, each packet receiver 96x, 96y has a clock... receive a corresponding one of the TNet Links 32. The CS FIFOs 102 operate to synchronize the incoming command/data symbols to the local clock of the packet receiver 96, buffering... 104x, coupled to the MUX 104y of the packet receiver 96y of the Y interface unit 24b by the cross-link connection 36(sub(x)). In similar fashion, information received at the Y port is coupled to the X interface unit 24a by the cross-link connection 36(sub(y)). In this manner, the command/data... packets received at one of the X, Y ports by the corresponding X, Y, interface unit 24a, 24b is passed to the other so that both will process and communicate the same information on to other components of the interface units 24 and/or memory 28.

Continuing with Fig. 6, depending upon which port X, Y... or the other of the CS FIFOs 102x, 102y for communication to the storage and processing logic 110 of the interface unit 24. The information contained in each 9-bit symbol is an 8-bit byte of... the encoding of which is discussed below with respect to Fig. 9. The storage and processing logic 110 will first translate the 9-bit symbols to 8-bit data or command... the outputs of the CS FIFOs 102x, 102y are also coupled to a command decode unit in addition to the MUX 104. The command decode unit operates to recognize command symbols (differentiating them from data symbols in a manner that is... below), decoding them to generate therefrom command signals that are applied to a receiver control unit, a state machine-based element that functions to control packet receiver operations.

As indicated above... at the output of the MUX 104, the receiver control portion of the storage control unit enables CRC check logic 106 to calculate a CRC symbol while the data symbols are... below, CS FIFOs are found not only in the packet receivers 96 of the interface units 24, but also at each receiving port of the routers 14 and the I/O... an even more important part, and perform a unique function, when a pair of sub-processor systems are operating in

duplex mode and the two CPUs 12A and 12B of the sub-processor systems 10A, 10B operate in synchronized, lock-step, executing the same instructions at the same time. When operating in this latter... ..difficult to ensure that the clocking regime of the routers 14A and 14B are exactly synchronized to those of the CPUs 12A and 12B - even when using frequency locked clocking. In... ..used to transmit symbols to a CPU 12 and the clock used by an interface unit 24 to receive those symbols.

The structure of the CS FIFO 102 is diagrammatically illustrated... ..i.e., a packet) or IDLE symbols - except during certain situations (e.g., reset, initialization, synchronization and others discussed below). As explained above, each symbol held in the transmit register 120...same symbol leaving the storage queue, allowing each symbol entering the storage queue 126 to settle before it is clocked out and passed to the storage and processing units 110x (and 110y) by the MUX 104x (and 104y). Since the transmit and receive clocks... ..functioning in duplex mode) operate to transmit symbols with near frequency clocking. Even so, clock synchronization FIFOs are used at these other ports to receive symbols transmitted with near frequency clocking, and the structure of these clock synchronization FIFOs are substantially the same as that used in frequency locked environments, i.e., that... ..of the storage queue 126 are nine bits wide; in near frequency environments, the clock synchronization FIFOs use symbol locations of the queue 126 that are 10 bits wide, the extra... ..the faster clock source. To handle this clock drift, the two pointers are effectively re-synchronized periodically.

When the CPUs 12 are paired and operating in duplex mode, all four interface

units 24 operate in lock-step to, among other things, transmit the same data and receive... ..simplex mode, each independent of the other, clocking need only be near frequency.

The interface unit 24 receives a SYNC CLK signal that is used in combination with a SYNC command symbol to initialize and synchronize the Rcv register 124 to the transmitting router 14. When using either near frequency or... ..102X preferably begin from some known state. Incoming symbols are examined by the storage and processing units 110 of the packet receivers 96. The storage and processing units look for, and act upon as appropriate, command symbols. Pertinent here is that when the... ..receives a SYNC command symbol it will be decoded and detected by the storage and processing unit 110. Detection of the SYNC command symbol by the storage and processing unit 110 causes assertion of a RESET signal. The RESET signal, under synchronous control of the SYNC CLK signal, is used to reset the input buffers (including the clock synchronization buffers) to predetermined states, and synchronize them to the routers 14.

The synchronization of the CS FIFOs 102 of the interface units 24 those ...one or both routers 14A, 14B is discussed more fully below in the section discussing synchronization.

Packet Transmitter:

Each interface unit 24 is assigned to transmit from and receive at only one of the X or Y ports of the CPU 12. When one of the interface units 24 transmits, the other operates to check the data being transmitted. This is an important... ..shows, in abbreviated form, the packet transmitters 94x, 94y of the X and Y interface units 24a, 24b, respectively. Both packet transmitters are identically constructed, so that discussion of one (packet... ..logic 152 that receives, from the BTE 88 or AVT 90 of the associated interface unit (here, the X interface unit 24a) the data to be transmitted - in doubleword (64-bit) format. The packet assembly logic... ..and Y ports: they are either symbols that make up a message packet in the process of being transmitted, or IDLE symbols, or other command symbols used to perform control functions... ..154, 156. The output of the multiplexer 154 connects to the X port. (The interface unit 24b connects the output of the multiplexer 154 to the Y port.) The multiplexer 156... ..sub(x)) to the checker logic 160 of the packet transmitter 94y (of the interface unit 24b).

A selection (S) input of the multiplexers receives a 1-bit output from an... ..is accessible to the MP 18 via an OLAP (not shown) formed in the interface unit 24, and is written with information that "personalizes," among other things,

the interface units 24 Here, the X/Y stage of the configuration register 162 configures the packet transmitter 94x of the X interface unit 24a to communicate the X encoder 150x output to the X port; the output of... ..traffic is present, the operation of the two packet interfaces 94 (and, thereby, the interface units 24 with which they are associated) are continually monitored. Should one of the checkers detect... ..will be asserted, resulting in an internal interrupt being posted for appropriate action by the processors 20.

Message packet traffic operates in the same manner. Assume, for the moment, that the... ..that information, a byte at a time, to the X encoder 150x of both interface units 96, which will translate each byte to encoded 9-bit form. The output of the... ..is checked with that from the packet transmitter 94x. Again, the operation of the interface units 24a, 24b, and the packet transmitters they contain, are inspected for error.

In the same... ..monitored.

Returning for the moment to Fig. 5, if the outgoing message packet is a processor initiated transaction (e.g., a read request), the processors 20 will expect a message packet to be returned in response. Thus, when the BTE... ..will issue a timeout signal to the interrupt logic (Fig. 14A) to thereby notify the processors 20 of the absence of a response to a particular transaction (e.g., a read... ..the access, to name just a few. Also, the area of memory of the memory unit 28 desired to be accessed are identified in the message packets by virtual or I... ..virtual addresses be translated to physical addresses of the memory 28. Finally, interrupts generated by units or elements external to the CPU 12A, are transmitted via message packets to interrupt the processors 20, which are also written to memory 28 when received. All ...this is handled by the interrupt logic and AVT logic 86, 90.

The AVT logic unit 90 utilizes a table (maintained by the processor 20 in memory 28) containing AVT entries for each possible external source permitted access to the memory 28. Each AVT entry identifies a specific source element or unit and the particular page (a page being nominally 4K (4096) bytes), or portion of a... ..expected" memory accesses. Expected memory accesses are those initiated by the CPU 12 (i.e., processors 20) such as a read request for information from an I/O device. These latter memory accesses are handled by a transaction sequence number (TSN) assigned to each processor initiated request. At about the time the read request is generated, the processors 20 will allocate an area of memory for the data expected to be received in... ..and 26b are, in turn, respectively coupled to the memory interfaces 70 of each interface unit 24a, 24b. The 64-bit doublewords are written to the memory 28 with the upper... ..check bits respectively from the memory interfaces 70 (70a, 70b) of each of the interface units 24a, 24b (Fig. 5).

Referring to Fig. 10, each memory interface 70 receives, from either the bus 82 from the processor interface 60 or the bus 83 from AVT logic 90 (see Fig. 5), of the associated interface unit 24, 64 bits of data to be written to memory. The busses 76 and 83... ..other for cross-checking between them. Thus, for example, the memory interface 70a (of interface unit 24a) will drive the MC 26a with the "upper" 32 bits of the 64 bits... ..are check bits, leaving 40 bits unused.

Access Validation:

As previously indicated, components of the processing system 10 external to the CPU 12A (e.g., devices of the I/O packet... ..not without qualification. Access validation, as implemented by the AVT logic 90 of the interface units 24, operates to prevent the content of the memory 28 from being ...Accesses to the memory 28 are validated by the AVT logic 90 of each interface unit 24 (Fig. 5), using all of six checks: (1) that the CRC of the message... ..also are permitted the particular message packet source.

The access validation mechanism of the interface unit 24a, AVT logic 88, is shown in greater detail in Fig. 11. Incoming message packets...and post an interrupt to the interrupt logic 86 (Fig. 5) for action by the processor 20.

The mask operation permits the size of the table of AVT entries to be varied. The content of the AVT mask register 175 is accessible to the processor 20, permitting the processors 20 to optionally select the size of the AVT entry table. A maximum AVT table... ..172 allows the AVT size to be matched to the needs of the system. A processing

system 10 that includes a larger number of external elements (e.g., the number of... amount of the memory space of memory 28 to the AVT entries. Conversely, a smaller processing system 10, with a smaller number of external elements will not have such a large... set to a logic "ZERO" indicate an nonexistent TNet address, outside the limits of the processing system 10. A received packet with a TNet address outside the allowable TNet range will... in Fig. 11 as being held in the AVT entry register 180 during the validation process. AVT entries have two basic formats: normal and interrupt. The format of a normal AVT... of the AVT input register 170) will result in an error being posted to the processor via an interrupt.

A 12-bit "Permissions" field is included in the AVT entry to... path=0). Denials are logged as interrupts with the interrupt logic, and reported to the processor 20 - if the E field is set to a state ("ONE") that enables error-reporting... e.g., to a "ONE"), the other fields (Upper Bound, etc.) gain new definitions for processing interrupt writes and managing interrupt queues. This is discussed in more detail below in connection... memory 28 will be handled. Set to one state, the requested write operation will be processed normally; set to a second state, write requests specifying addresses with a fractional cache line... be written to a specific queue (interrupt queue) in memory 28, with signalling provided the processors 20 to indicate that an interrupt has been received and "posted," and ready for servicing by the processors 20. Since the interrupt queues are at specific memory locations, the processor can obtain the interrupt data when needed.

An AVT interrupt entry for an interrupt may... by the interrupt logic 86, and extracted from the head of the queue by the processor 20 when servicing the interrupt.

The AVT interrupt entry also includes a 20-bit segment ("Source ID") containing source ID information, identifying the external unit seeking attention by the interrupt process. If the source ID information of the AVT interrupt entry does not match that contained... class" of the interrupt that is used to determine the interrupt level set in the processor 20 (described more fully below); (2) a queue number that is used to select, as... capability to deliver interrupts to a CPU 12 for servicing. For example, an I/O unit may be unable to complete a read or write transaction issued by a CPU because... identify the recipient. These and other errors, exceptions, and irregularities, noted by the I/O units, or the I/O Interface elements, can become the a condition that requires the intervention... the AVT entry register 180 for use by the interrupt logic 86 of the interface unit 24 (Fig. 5), illustrated in greater detail in Fig. 14A.

It is interrupt logic 86... four circular queues specified by the base address information contained in the AVT entry. The processor(s) 20 will then be notified, and it will be up to them as to... selected tail queue register 256 by combiner circuit 270, the output of which is the processed by the "mod z" circuit 273 to turn new offset into the queue at which... signal. The Queue Full warning signal becomes an "intrinsic" interrupt that is conveyed to the processor units 20 as a warning that if the matter is not promptly handled, later-received interrupt... will be discarded.

Incoming message packet interrupts will cause interrupts to be posted to the processor 20 by first setting one of a number of bit positions of an interrupt register 280. Multi-entry queued interrupts are set in interrupt registers 280a for posting to the processor 20; single-entry queue interrupts use interrupt register 280b. Which bit is set depends upon... multi-entry queued interrupts, soon after a multi-entry queued interrupt is determined, the interface unit will assert a corresponding interrupt signal (I1) that is applied to decode circuit 283. Decode... of register 280a to set, thereby providing advance information concerning the received interrupt to the processor(s) 20, i.e., (1) the type of interrupt posted, and (2) the class of... to one another by a compare circuit 279. The update register is writable by the processor 20 to select a register pair for comparison. If the content of the two selected... cleared.

Digressing for the moment, there are two basic types of interrupts that concern the processors 20: those interrupts that are communicated to the CPU 12 by message packets, and those... the seven interrupt postings to a latch 288,

from which they are coupled to the processor 20 (20a,20b) which has an interrupt register for receiving holding the postings.

In addition... ..change in interrupts (either an interrupt has been serviced, and its posting deleted by the processor 20, or a new interrupt has been posted), a "CHANGE" signal will be issued to the processor interface 60 to inform it that an interrupt posting change has occurred, and that it should communicate the change to the processor 20.

Preferably, the AVT entry register 180 is configured to operate like a single line... ..such as set-associative, fully-associate, or direct-mapped, to name a few.

Coherency:

Data processing systems that use cache memory have long recognized the problem of coherency: making sure that... ..the incoming packet is permitted access are applied to a boundary crossing (Bdry Xing) check unit 219. Boundary check unit 219 also receives an indication of the size of the cache block the CPU 12... ..Len field of the header information from the AVT input register 170. The Bdry Xing unit determines if the data of the incoming packet is not aligned on a cache boundary...time an interrupt will be written to the queued interrupt register 280, to alert the processors 20 that a portion of the incoming data is located in the special queue.

In... ..not, the packet (both header and data) is written to a special queue, and the processors so notified by the intrinsic interrupt process described above. The processors may then move the data from the special queue to cache 22, and later write... ..the cache 22 and the memory 28 is preserved.

Block Transfer Engine (BTE):

Since the processor 20 is inhibited from directly communicating (i.e., sending) information to elements external to the... ..indirect method of information transmission.

The BTE 88 is the mechanism used to implement all processor initiated I/O traffic to transfer blocks of information. The BTE 88 allows creation of... ..BTE registers 300, 302 whose content is coupled to the MUX 306 (of the interface unit 24a; Fig. 5) and used to access the system memory 28 via the memory controllers... ..BTE data structure 304 in the memory 28 of the CPU 12A (Fig. 2). The processors 20 will write a data structure 304 to the memory 28 each time information is... ..begin on a quadword boundary, and the BTE registers 300, 302 are writable by the processors 20 only. When a processor does write one of the BTE registers 300, 302, it does so with a word... ..the request bit (rc0, rc1) to a clear state, which operates to initiate the BTE process, which is controlled by the BTE state machine 307.

The BTE registers 300, 302 also... ..cause (ec) bit differentiates time-outs and NAKs.

When information is being transferred by the processors 20 to an external unit, the data buffer portion 304b of the data structure 304 holds the information to be transferred. When information from an external unit is received by the processors 20, the data buffer portion 304b is the location targeted to hold the read response information.

The beginning of the data structure 304, portion 304a written by the processor 20, includes an information field (Dest), identifying the external element which will receive the packet... ..the transmitted data is to be written. This information is used by the packet transmitter unit 120 (Fig. 5) to assemble the packet in the form shown in Figs. 3-4...list (el) bit, when set, indicates the end of the chain, and halts the BTE processing.

The interrupt completion (ic) bit, when set, will cause the interface unit 24a to assert an interrupt (BTEComp) which sets a bit in the interrupt register 280... ..the chain pointer).

The interrupt time-out (it) bit, when set, will cause the interface unit 24a to assert an interrupt signal for the processor 20 if the acknowledgement of the access times-out (i.e., if the request timer... ..time), or elicits a NAK response (indicating that the target of the request could not process the request).

Finally, if the check sum (cs) bit is set, the data to be... ..containing the data from which the check sum was formed.

To sum up, when the processors 20 of the CPU 12A desire to send data to an external unit, they will write a data structure 304 to the memory 28, comprising identifier information in portion 304a of the data structure, and the data in the buffer portion 304b. The processors 20 will then determine the priority of the data and will write the BTE register... ..information, and sent.

If the data structure 304 indicates a read request (i.e., the processors 20 are seeking data from an external unit - either an I/O device or a CPU 12), the Len and Local Buffer Ptr... ..receiver 100 (Fig. 5) until the local memory write operation is executed.

Responses to a processor-generated read request to an external unit are not processed by the AVT table logic 146. Rather, when the processors 20 set up the BTE data structure, a transaction sequence number (TSN) is assigned the... ..the BTE 88, which will be an HAC type packet (Fig. 4) discussed above. The processors 20 will also include an memory address in the BTE data structure at which the...302, assume that the foregoing transfer of data from the CPU 12A to an external unit is of a large block of information. Accordingly, a number of data structures would be set up in memory 28 by the processors 20, each (except the last) including a chain pointer to additional data structures, the sum... ..sent. Assume now that a higher priority request is desired to be made by the processors 20. In such a case, the associated data structure 304 for such higher priority request... ..with another BTE operation descriptor.

Memory Controller:

Returning, for the moment, to Fig. 2, interface units 24a, 24b access the memory 28 via a pair of memory controllers (MC) 26a, 26b. The Mcs provide a fail-fast interface between the interface units 24 and the memory 28. The Mcs 26 provide the control logic necessary for accessing... ..in dynamic random access memory (DRAM) logic). The Mcs receive memory requests from the interface units 24, and execute reads and writes as well as providing refresh signals to the DRAMs... ..to provide a 72 bit data path between the memory array 28 and the interface units 24a, 24b, which utilize an SBC-DBD-SbD ECC scheme, where $b=4$, on a... ..26a, 26b to work together and simultaneously supply a 64-bit word to the interface units 24 with minimum latency, one-half of which (D0) comes from the MC 26a, and the other half (D1) comes from the other MC 26b. The interface unit 24 generate and check the ECC check bits. The ECC scheme used will not only... ..26 bus 25, as well as in internal registers.

From the viewpoint of the interface units 24, the memory 28 is accessed with two instructions: a "read N doubleword" and a... ..doubleword read or a block read format. The signal called "data valid" tells the interface units 24 two cycles ahead of time that read data is being returned or not being returned.

As indicated above, the maintenance processor (MP 18; Fig. 1A) has two means of access to the CPUs 12. One is... ..18 will write a register contained in the OLAP 285 with instructions that permit the processors 20 to build an image of a sequence of instructions in the memory that will permit them (the processors 20) to ...to transfer instructions and data from an external (storage) device that will complete the boot process.

The OLAP 285 is also used by the processors 20 to communicate to the MP 18 error indications. For example, if one of the interface units 24 detect a parity error in data received from the memory controller 26, it will... ..and address transfers on the bus 25 between the MC 26a and the corresponding interface unit 24a. The addressing and data transfers on the DRAM data bus, as well as generation... ..the CPU 12.

Packet Routing:

The message packets communicated between the various elements of the processing system 10 (e.g., CPUs 12A, 12B, and devices coupled to the I/O packet... ..First, each TNet Link L connects to an element (e.g., router 14A) of the processing system 10 via a port that has both receive and transmit capability. Each transmit port... ..cycle (i.e., each clock period) of the T(underscore)CFk so that the clock

synchronization FIFO at the receiving end of the transmission will maintain synchronization.

Clock synchronization is dependent upon the mode in which the processing system 10 is operated. If operating in the simplex mode in which the CPUs 12A... connect directly to the CPUs may drift with respect to each other. Conversely, when the processing system 10 operates in a duplex mode (e.g., the CPUs operate in synchronized, lock-step operation), the clocks between routers 14 and the CPUs 12 to which they... not necessarily phase-locked).

The flow of data packets between the various elements of the processing system 10 is controlled by command symbols, which may appear at any time, even within... initiated by a CPU 12, or MP 18, and promulgated to all elements of the processing system 10 by the routers 14 to communicate an event requiring software action by all... command symbol is used in conjunction with near frequency operation as an aid to maintaining synchronization between the two clock signals that (1) transfer each symbol to, and load it in each receiving clock synchronization FIFO, and (2) that retrieves symbols from the FIFO.

SLEEP: This command symbol is sent by any element of the processing system 10 to indicate that no additional packet (after the one currently being transmitted, if... received.

SOFT RESET (SRST): The SRST command symbol is used as a trigger during the processes ("synchronization" and "reintegration," described below) that are used to synchronize symbol transfers between the CPUs 12 and the routers 14A, 14B, and then to place... SYNC command symbol is sent by a router 14 to the CPU 12 of the processing system 10 (i.e., the sub-processor systems 10A/10B) to establish frequency-lock synchronization between CPUs 12 and routers 14A, 14B prior to entering duplex mode, or when in duplex mode to request synchronization, as will be discussed more fully below. The SYNC command symbol is used in conjunction... or duplex to simplex), among other things, as discussed further below in the section on Synchronization and Reintegration.

THIS LINK BAD (TLB): When any system element receiving a symbol from a TNet link L (e.g., a router, a CPU, or an I/O unit) notes an error when receiving a command symbol or packet, it will send a TLB... identical pairs of symbols that are compared to one another when pulled from the clock synchronization FIFOs. The DVRG command symbol signals the CPU 12 that a mis-compare has been noted. When received by the CPUs, a divergence detection process is entered whereby a determination is made by the CPUs which CPU may be failing... command symbols described above operate to control message flow between the various elements of the processing system 10 (e.g., CPUs 12, router 14, and the like), using principally the BUSY however, an "end node" (i.e., a CPU 12 or I/O unit 17 - Fig. 1) may not assert backpressure because one of its transmit ports is backpressured... Improperly addressed packets are discarded by the router 14.

When a system element of the processing system 10 receives a BUSY command symbol on a TNet link L on which it... other command symbols (READY, BUSY, etc.).

Whenever a TNet port of an element of the processing system 10 detects receipt of a READY command symbol, it will terminate transmission of FILL... receives.

As will be seen, all elements (e.g., router 14, CPUs 12) of the processing system 10 that connect to a TNet link L for receiving transmitted symbols will receive those symbols via a clock synchronization (CS) FIFO. For example, as discussed above, the interface units 24 of CPUs 12 include all CS FIFOs 102x, 102y (illustrated in Fig. 6). The... depth to allow for speed matching, and the elastic FIFOs must provide sufficient depth for processing delays that may occur between transmission of a BUSY command symbol during receipt of a... another data byte in packet B. As packet A progresses to the next router, the process would be repeated. If the router 14 displaces more data bytes than the FIFO can... irrespective of its own findings.

SLEEP Protocol:

The SLEEP protocol is initiated by a maintenance processor via a maintenance interface (an on-line access port - OLAP), described below. The SLEEP protocol... ..reintegrate a slice of the system 10. Routers 14 must be idle (no packets in process) in order to change modes without causing data loss or corruption. When a SLEEP command symbol is received, the receiving element of processing system 10 inhibits initiation of transmission of any new packet on the associated transmit port... ..The HALT command symbol provides a mechanism for quickly informing all CPUs 12 in a processing system 10 that is necessary to terminate I/O activity (i.e., message transmissions between... ..CPUs that receive HALT command symbols on either of their receive ports (of the interface units 24) will post an interrupt to the interrupt register 280 if the system halt interrupt... ..interrupt; Fig. 14A).

The CPUs 12 may be provided with the ability to disable HALT processing. Thus, for example, the configuration registers 75 of the interface units 24 can include a "halt enable register" that, when set to a predetermined state (eg., ZERO) disables HALT processing, but reporting detection of a HALT symbol as an error.

Router Architecture:

Referring now to... ..simplified block diagram of the router 14A is illustrated. The other routers 14 of the processing system 10 (e.g., routers 14B, 14', etc.) are of substantially identical construction and, therefore...these ports 4, 5 are structured to operate in a frequency locked environment when a processing system 10 is set for duplex mode operation. In addition, when in duplex mode, a... ..5)) will receive the command/data symbols from the CPUs, pass them through the clock synchronization FIFOs 518 (discussed further below), and compare each symbol exiting the clock synchronization FIFOs with a gated compare circuit 517. When duplex operation is entered, a configuration register... ..517 to activate the symbol by symbol comparison of the symbols emanating from the two synchronization FIFOs 518 of the router input logic 502 for the ports 4 and 5. Of... ..to that received, at substantially the same time, by the other port input.

To maintain synchronization in the duplex mode, the two port outputs of the router 14A that transmit to... ..mode, are duplicated by the routers 14, and returned to both CPUs.) The output logic units 504(sub(4)), 504(sub(5)) that are coupled directly to the CPUs 12 will... ..message packet identifies only one of the duplexed CPUs 12, e.g., CPU 12A) in synchronized fashion, presenting those symbols in substantially simultaneous fashion to the two CPUs 12. Of course, the CPUs 12 (more accurately, the associated interface units 24) receive the transmitted symbols with synchronizing FIFOs of substantially the same structure as that illustrated in Fig. 7A so that, even... ..from the FIFO structures by both CPUs 12 on the same instruction cycle, maintaining the synchronized, lock-step operation of the CPUs 12 required by the duplex operating mode.

As will... ..conjunction with configuration data written to registers contained in control logic 509 by the maintenance processor 18 (via the on-line access port 285' and serial bus 19A; see Fig. 1A...links L. The input logic 505 of each port input 502 also assists in maintaining synchronization - at least for those ports sending symbols in the near-frequency environment - by removing received... ..slower-receiving element receiving symbols from a faster-sending element could overload the input clock synchronization FIFO of the slower-receiving element. That is, if a slower clock is used to pull symbols from the clock synchronization FIFO put there by a faster clock, ultimately the clock synchronization FIFO will overflow.

The preferred technique employed here is to periodically insert SKIP symbols in... ..stream to avoid, or at least minimize, the possibility of an overflow of the clock synchronization FIFO (i.e., clock synchronization FIFO 518; Fig. 20A) of a router 14 (or CPU 12) due to a T... ..being slightly higher in frequency than the local clock used to pull symbols from the synchronization FIFO. Using SKIP symbols to by-pass a push (onto the FIFO) operation has the... ..stall each time a SKIP command symbol is received so that, insofar as the clock synchronization FIFO is concerned, the transmitting clock that accompanied the SKIP symbol was missing.

Thus, logic... ..the port inputs 502 will recognize, and key off receipt of, SKIP command symbols for synchronization in the near frequency clocking environment so that nothing is pushed onto the FIFO, but... ..14, or between routers 14, or between a router 14 and an I/O interface unit 16A - Fig. 1) at a 50 Mhz rate, this allows for a

worst case frequency... symbol by supplying FILL or IDLE symbols (which are received and pushed onto the clock synchronization FIFOs, but are not passed to the elastic FIFOs). In short, each elastic FIFO 506...received symbols are then communicated from the input register 516 and applied to a clock synchronization FIFO 518, also by the T(underscore)Clk. The clock synchronization FIFO 518 is logically the same as that illustrated in Figs. 8A and 8B, used in the interface units 24 of the CPUs 12. Here, as Fig. 20A shows, the clock synchronization FIFO 518 comprises a plurality of registers 520 that receive, in parallel, the output of... 516. Associated with each of the registers 520 is a two-stage validity (V) bit synchronizer 522, shown in greater detail in Fig. 20B, and discussed below. The content of each... registers 520, together with the one-bit content of each associated two-stage validity bit synchronizer 522, are applied to a multiplexer 524, and the selected register/synchronizer pulled from the FIFO, and coupled to the elastic FIFO 506 by a pair of... is determined the state of the Push Select signal provided by a push pointer logic

unit 530; and, selection of which register 520 will supply its content, via the MUX 524... and loading of the register 520 selected by the push pointer logic 530. Similarly, the synchronization FIFO control logic 534 receives the clock signal local to the router (Rcv Clk) to... pointer logic 532.

Digressing for a moment, and referring to Fig. 20B, the validity bit synchronizer 522 is shown in greater detail as including a D-type flip-flop 541 with... 530 (Fig. 20A) selects the register 520 of the FIFO with which the validity bit synchronizer is associated for receipt of the next symbol - if not a SKIP symbol.

The delay... Truth Table, below). The D-type flip-flop 543 acts as an additional stage of synchronization, ensuring a stable level at the V output relative to the local Rec Clk. The... flip-flop 542, allowing the Pull signal (a periodic pulse from the sync FIFO Control unit 534) to clear the validity bit on this validity synchronizer 522 when the associated register 520 has been read. (Table omitted)

In summary, the validity synchronizer 522 operates to assert a "valid" (V) signal when a symbol is loaded in a...blocked from being routed out a particular port because another message is already in the process of being routed out that port. However, that other message in turn is also blocked...an incoming message packet bound for the CPUs will be replicated by the crossbar logic unit by routing the message packet to both port output 504(sub(4)) and 504(sub... P) identifies which of path (X or Y) should be used for accessing two sub-processing the device.

The routers 14 provide a capability of constructing a large, versatile routing network for, for example, massively parallel processing architectures. Routers are configured according to their location (i.e., level) in the network by...j)) and 509(sub(k)) are such that bits "def" are used in the algorithmic process, then bits "abc" of the Region ID are compared to the content of the Device... the route to default register 509(sub(f)) to the final stage of the selection process: check logic 602. Check logic 602 operates to check the status of the port output...a lower level router, and may be located in one or another of the sub-processing systems 10A, 10B. Whether a router is an upper level or lower level router depends... of CPUs 12 and I/O devices 16 to one another, forming a massively parallel processing (MPP) system. Other such MPP systems may exist, and it is those routers configured as... captured. As soon as the message packet's Destination ID is so captured, the selection process begins, proceeding to the development of a target port address that will be used to...an error that will be posted to the MP18 via the router's (or interface unit's) OLAP for action.

Digressing, it should be appreciated that these protocol rules observed by the routers 14 are also observed by the CPUs 12 (i.e., interface units 24) and I/O packet interfaces 17.

Finally, when the router 14A is in the... directly with the CPUs 12A, 12B, and duplex mode is used, a duplex operation logic unit 638 is utilized to coordinate the port output connected to one of the CPUs 12A... was able to write instructions to the OLAP 285 that would be executed by the processors 20 to build a small memory image and routine to permit the CPU 12 to... the clock generation circuit design. There will be one clock generator circuit in

each sub-processor system 10A/10B (Fig. 1) to maintain synchronism. Designated generally with the reference numeral 650... used by the various elements (e.g. CPU 12, routers 14, etc.) of the sub-processor system containing the clock circuit 650 (e.g., 10A).

The clock generator 654 is shown... The 50 Mhz clock signals produced by the counter 663 are distributed throughout the sub-processor system where needed.

Turning now to Fig. 25, there is illustrated the interconnection and use... the clock circuits 650 used to develop synchronous clock signals for a pair of sub-processor systems 10A, 10B (Fig. 1) for frequency locked operation. As illustrated in Fig. 25, the two CPUs 12A and 12B of the sub-processor systems 10A, 10B each have a clock circuit 650, shown in Fig. 25 as clock... 654B of both CPUs 12. A driver and signal line 667 interconnects the two sub-processor systems to deliver the M(underscore)CLK signal developed by the oscillator circuit 652A to the clock generator 654B of the sub-processor system 10B. For fault isolation, and to maintain signal quality, the M(underscore)CLK signal is delivered to the clock generator 654A of the sub-processor system 10A through a separate driver and a loopback connection 668. The reason for the... the cable (not shown) will establish the connection shown in Fig. 25 between the sub-processor systems 10A, 10B; connected another way, the connections will be similar, but the oscillator 652B... Fig. 25, the M(underscore)CLK signal produced by the oscillator circuit 652A of sub-processing system 10A is used by both sub-processing systems 10A, 10B as their respective SYNC CLK signals and the various other clock signals... produced by the clock generators 654A, 654B. Thereby, the clock signals of the paired sub-processing systems 10A, 10B are synchronized for the frequency locked operation necessary for duplex mode.

The VCXOs 662 of the clock... This allows both clock generators 654A, 654B to continue to provide to the two sub-processing systems 10A, 10B clock signals in the face of improper operation of the oscillator circuit 652A, although the sub-processor systems may no longer be frequency-locked.

The LOCK signals asserted by the phase comparators... LOCK signal signifies that the 50 Mhz signals produced by a clock generator 654 are synchronized, both in phase and in frequency, to the M(underscore)CLK signal. Thus, if either... signal that accompanies the symbol stream, and is used to push symbols onto the clock synchronizing FIFO of the receiving element (router 14, or CPU 12) is substantially identical in frequency... not phase, to that of the receiving element used to pull symbols from the clock synchronization FIFOs. For example, referring to Fig. 23, which illustrates symbols being sent from the router... clock (Local Clk). The former (Rev Clk) is used to push symbols onto the clock synchronization FIFOs 126 of each CPU, whereas the latter is used to pull symbols from the... much higher frequency clock signal. In such situations provision must be made to ensure that synchronization is maintained between the two CPUs as to symbols pulled from the clock synchronization FIFOs 126 of each.

Here, a constant ratio clocking mechanism is used to control operation of the two clock synchronization FIFOs 126, providing the clock signal that pulls symbols from the two FIFOs at the... control mechanism is shown, designated with the reference numeral 70. As Fig. 26A illustrates, clock synchronization FIFO control mechanism 700 includes an pre-settable, multi-stage serial shift register 702, the... ratio of the clock signal at which symbols are communicated and pushed onto the clock synchronization FIFOs 126 to the frequency of the clock signal used locally. Here, a 15 stages... that will be used as the Local Clk signal to pull symbols from the clock synchronization FIFOs 126, and to operate (update) the pull pointer counter 130. The selected output is... of the CPU 12 to the clock signal used to push symbols onto the clock synchronization FIFO 126, Rev Clk, the serial shift register is preset so that M stages of... duplexed CPUs 12 with a 50 Mhz clock. Thus, symbols are pushed onto the clock synchronization FIFOs 126 of the CPUs at a 50 Mhz rate. Assume further that the clock... of the MUX 704, which produces the clock signal that pulls symbols from the clock synchronization FIFOs 126, Rev Clk, will contain, for each 100 ns period, five clock pulses. Thus... five symbols will be pushed onto, and five symbols will be pulled from, the clock synchronization FIFOs 126.

This example is symbolically shown in Fig. 26B, while the timing diagram shown... ..labelled "IN" in Fig. 27) of the Rev Clk will push symbols onto the clock synchronization FIFOs 126. During that same 100 ns period, the serial shift register 702 circulates a... ..clocks which would require additional storage (i.e., an increase in the size of the synchronization FIFO) and impose more latency.

The constant ratio clock circuit presented here (Figs. 26) is... ..frequency to a clock regime of a different, higher frequency. The use of a clock synchronization FIFO is necessary here for compensating effects of signal delays when operating in synchronized, duplexed mode to receive pairs of identical command/data symbols from two different sources. However...so long as there are at least two registers in the place of the clock synchronization FIFO. Transferring data from a higher-frequency clock regime to a lower frequency clock regime... ..a wide range of possible clock ratios.

I/O Packet Interface:

Each of the sub-processor systems 10A, 10B, etc. will have some input/output capability, implemented with various peripheral units, although it is conceivable that the I/O of other sub-processor systems would be available so that a sub-processing system may not necessarily have local I/O. In any event, if local I/O... ..device (e.g., a signal line) would be received by the I/O packet interface unit 16 and used to form an interrupt packet that is sent to the CPU 12... ..OLAP bus, configuration information.

On-Line Access Port:

The MP 18 connects to the interface unit 24, memory controller (MC) 26, routers 14, and I/O packet interfaces with interface signals... ..OLAP 258 is essentially the same, regardless of what element (e.g. router 14, interface unit 24, etc.) it is used with. Fig. 28 diagrammatically illustrates the general structure of the... ..circuit chip used to implement certain of the elements discussed herein. For example, each interface

unit 24, memory controller 26, and router 14 is implemented by an application specific integrated circuit... ..of the OLAP 158 shown in Fig. 28 describes the OLAP associated with the interface unit 24, the MC 26, and the router 14 of the system.

As Fig. 28 shows...asymmetric variables, a "soft-vote" (SV) logic element 900 (Fig. 30A) is provided each interface unit 24 of each CPU 12. As Fig. 30 illustrates, the SV logic elements 900 of each interface unit 24 are connected to one another by a 2-bit SV bus 902, comprising bus lines 902a and 902b. Bus lines 902a carry one-bit values from the interface units 24 of CPU 12A to those of CPU 12B. Conversely, bus line 902b carries one... ..the CPU 12A.

Illustrated in Fig. 30B, is the SV logic element 900a of interface unit 24a of CPU 12A. Each SV logic element 900 is substantially identical in construction and... ..900a should be understood as applying equally to the other logic elements 900a (of interface unit 24b, CPU 12A), and 900b (of the interface units 24a, 24b of CPU 12B) unless noted otherwise. As Fig. 30B illustrates, the SV logic... ..the logic elements 900a (as well as its own). In this manner the two interface units 24a, 24b of the CPU 12A can communicate asymmetrical variables to each other.

In a... ..to the remote register 907 of logic element 902a (and that of the other interface unit 24b).

The logic elements 902 form a part of the configuration registers 74 (Fig. 5). Thus, they may be written by the processor unit(s) 20 by communicating the necessary data/address information over at least a portion of... ..local and remote registers 906 and 907.

The MUX 914 operates to provide each interface unit 24 of CPU 12A with selective use of the bus line 902a for the SV logic elements 900a, or for communicating a BUS ERROR signal if encountered during the reintegration process (described below) used to bring a pair of CPUs 12 into lock-step, duplex operation... ..same time, write the enable registers 912 of the logic element 900 of both interface units 24 of each CPU. One of the two logic elements

900 of each CPU will... it is the output enable registers 912 associated with the logic elements 900 of interface units 24a of both CPUs 12A, 12B that are written to enable the associated drivers 916. Thus, the output registers 904 of the interface units 24a of each CPU will be communicated to the bus lines 902; that is, the... to the bus line 902a, while the output register associated with logic element 900b, interface unit 24a of CPU 12B is communicated to bus line 902b. The CPUs 12 will both... again written by each CPU, followed again by reading the remote input registers 907. This process is repeated, one bit at a time, until the entire variable is communicated from the... each CPU 12 to the remote input register of the other. Note that both interface units 24 of CPU 12B will receive the bit of asymmetric information.

One example of use... elements 900 are also used to communicate bus errors that may occur during the reintegration process to be described. When reintegration is being conducted, a REINT signal will be asserted. As... ERROR signal is selected by the MUX 914 and communicated to the bus line 902a.

Synchronization:

Proper operation of the sub-processing systems 10A, 10B (Figs. 1A, 2) whether operating independently (simplex mode), or paired and operating in synchronized lock-step (duplex mode), requires assurance that data communicated between the CPUs 12A, 12B and... routers 14A, 14B will be received properly, and that any initial content of the clock synchronization FIFOs 102 (of CPUs 12A, 12B; Fig. 5) and 519 (of routers 14A, 14B; Fig... erroneously interpreted as data or commands. The push and pull pointers of the various clock synchronization FIFOs 102 (in the CPUs 12) and 518 (in the routers 14) need to be... apart, and presetting the associated FIFO queues to some known state. This done, all clock synchronization FIFOs are initialized for near frequency operation. ...in order to properly implement the lock-step operation of duplex mode operation, the clock synchronization FIFOs must be synchronized to operate with the particular source from which they receive data in order accommodate any... 14A, 14B to the CPUs 12A, 12B must be accounted for. It is the clock synchronization FIFOs 102 of the paired CPUs 12 that operate to receive message packet symbols, adjust... and present symbols to the two CPUs in a simultaneous manner to maintain lock-step synchronization necessary for duplex mode operation.

In similar fashion, each symbol received by the routers 14A... the CPUs (which is discussed further hereinafter). Again, it is the function of the clock synchronization FIFOs 518 of the routers 14A, 14B that receive message packets from the CPUs 12... so that the symbols received from the two CPUs 12 are retrieved from the clock synchronization FIFOs simultaneously.

Before discussing how the clock synchronization FIFOs of the CPUs and routers are reset, initialized, and synchronized, an understanding of their operation to maintain synchronous lock-step duplex mode operation is believed helpful. Thus, referring for the moment to Fig. 23, the clock synchronization FIFOs 102 of the CPUs 12A, 12B that receive data, for example, from the router... T Clk, from the router 14A to the CPU 12B.

Consider operation of the clock synchronization FIFOs 102(sub(x)), 102(sub(y)), to receive identical symbol streams during duplex operation... held by the push and pull pointer counters 128, 130 for the CPU 12A (interface unit 24a), and the content of each of the four storage locations (byte 0. ..., byte 3... 6 show the same thing for the FIFO 102(sub(y)) of CPU 12B interface unit 24a for each symbol of the duplicated symbol stream.

Assuming the delay 640 is no... locations of the queues 126. This is because (1) the FIFOs 102 have been synchronized to operate in synchronism (a process described below), and (2) the push pointer counters 128 are clocked by the clock signal... of the symbol stream transmitted by the router 14A will be pulled from the clock synchronization FIFOs 102 of the CPUs 12A, 12B simultaneously, maintaining the required synchronization of received data when operating in duplex mode. In effect, the depths of the queues... order to achieve the operation just described with reference to Table 6, the reset and synchronization process shown in Fig 31A is used. The process not only initializes the clock synchronization FIFOs 102 of the CPUs 12A, 12B for duplex mode operation, but also operates to adjust the clock synchronization FIFOs 518 (Fig. 19A) of the CPU ports of each of

the routers 14A, 14B for duplex operation. The reset and synchronization process uses the SYNC command symbol to initiate a time period, delineated by the SYNC CLK signal 970 (Fig. 31B), to reset and initialize the respective clock synchronization FIFOs of the CPUs 12A and 12B and routers 14A, 14B. (The SYNC CLK signal... It is of a lower frequency than that used to receive symbols by the clock synchronization FIFOs, T(underscore)Clk. For example, where T(underscore)Clk is approximately 50 MHz, the... signal is approximately 3.125 MHz.)

Turning now to Fig. 31A, the reset and initialization process begins at step 950 by switching the clock signals used by the CPUs 12A, 12B and routers 14A, 14B as the transmit (T(underscore)Clk) and the unit's local clock (Local Clk) clock signals so that they are derived from the same... In addition, configuration registers in the CPUs 12A, 12B (configuration registers 74 in the interface units 24) and the routers 14A, 14B (contained in control logic unit 509 of routers 14A, 14B) are set to the FreqLock state.

The following discussion involves step 952, and makes reference to the interface unit 24 (Fig.5), router 14A (Fig. 19A) and Figs. 31A and 31B. With the clock... otherwise be sent followed by a self-addressed message packet. Any message packet in the process of being received and retransmitted when the SLEEP command symbols are received and recognized by... per the destination address). The SLEEP command symbol operates to "quiece" router 14A for the synchronization process. The self-addressed message packet sent by the CPU 12A, when received back by the... message packet sent after the SLEEP command symbol would necessarily have to be the last processed by the router 14A.

At step 954 the CPU 12A checks to see if it...the router will assert a RESET signal 972 that is applied to the two clock synchronization FIFOs 518 contained in the input logic 505(sub(4)), 505(sub(5)) of the... receive symbols directly from CPUs 12A, 12B. RESET, while asserted, will hold the two clock synchronization FIFOs 518 in a temporarily non-operating reset state with the push and pull pointer... As each of the CPUs 12 receive SYNC symbols are detected by the storage and processing units of the packet receivers 96 (Figs. 5 and 6) cause the RESET signal to be asserted by the packet receivers 96 (actually, storage and processing elements 110; Fig. 6) of each CPU 12. the RESET signal is applied to the... 4)), CPUs 12 and routers 14A, 14B de-assert the RESET signals, and the clock synchronization FIFOs of the CPUs 12A, 12, and routers 14A, 14B are released from their reset... the delay, the router 14A and CPUs 12 resume pulling data from their respective clock synchronization FIFOs and resume normal operation. The clock synchronization FIFOs of the router 14A begin pulling symbols from the queue (previously set by RESET... from the CPU 12A with the T(underscore)Clk will be pushed onto the clock synchronization FIFO at, for example, queue location 0 (or whatever other location pointed to by the... 0 (or whatever other location the push pointer was set to by RESET). The clock synchronization FIFOs of the router 14A are now synchronized to accommodate whatever delay 640 may be present in one communications path, relative to the... and the CPUs 12A, 12B.

Similarly, at the same virtual time, operation of the clock synchronization FIFOs 102 of both CPUs 12A, 12B is resumed, synchronizing them to the router 14A. Also, the CPUs 12A, 12B quit sending the SLEEP command... in favor of READY symbols, and resume message packet transmission, as appropriate.

That completes the synchronization process for the router 14A. However, the process must also be performed for the router 14B. Thus, the CPU 12A returns to step... however, assuming that the CPUs 12A, 12B are operating in duplex mode, the method and apparatus used to detect and handle a possible error, resulting in divergence of the CPUs from... via a message packet destined for a peripheral device of one or the other sub-processor systems 10A, 10B. Depending upon the destination of the outgoing message packet, step 1002 will... router 14 will issue an ERROR signal to the router control logic 509, causing the process to move to step 1004 where the router 14 detecting divergence will transmit a DVRG... time outs to occur. A router detecting divergence (without also detecting any simple link error) buys itself time to check the CRC of the received message packet by waiting for the... router 14, or received, all further message packets received from the CPUs and in the process of being routed

when divergence was detected, or the DVRG symbol received, will be passed... 1010) contained in one of the configuration registers 74 (Fig. 5) of the interface unit 24 of each CPU.

Returning for the moment to step 1006, the determination of which... "local" is meant to refer to the router 14A, 14B contained in the same sub-processor system 10A, 10B as the CPU. For example, referring to Fig. 1A, router 14A is... bit mentioned above: the bit contained in one of the configuration registers 74 of interface unit 24 (Fig. 5) of each CPU. When set to a first state, that particular CPU... the other CPU. In response, the state machines (not shown) within the control and status unit 509 (Fig. 19A) changes the "favorite" bits described above.

A few examples may facilitate understanding... DVRG symbol will echo that symbol to the routers 14A, 14B, start its internal divergence process timer, and begin determination of whether to continue or terminate. Having received a TLB symbol... to diverge with no errors reported. This can happen only if software (running on the processors 20) uses known divergent data to alter state. For example, suppose each CPU 12 has... number of the CPU 12A will differ from that of the CPU 12B. If the processors use the serial number to change the sequence of instructions executed (say, by branching if the serial number comes after some value) or to modify the value contained in a processor register, the complete "state" of the CPUs 12 will differ. In such cases, the "asymmetrical... of the primary CPU simply allows one CPU, and thereby the system 10, to continue processing without software intervention.

- An error at the output of the interface unit 24 of a CPU 12 will be detected by the router 14A, 14B, depending upon... router 14A, 14B that connects to a CPU 12 will be detected by the interface unit 24 of the affected CPU. The CPU will send a TLB symbol to the faulty... possible failure and, without external intervention, and transparently to the system user, remove the failing unit (CPU 12A or 12B, or router 14A or 14B) from the system to obviate or... reintegration." The discussion will refer to the CPUs 12A, 12B, routers 14A, 14B, and maintenance processor 18A, 18B shown forming parts of the processing system 10 illustrated in Fig. 1A. In addition, discussion will refer to the processors 20a, 20b, the interface units 24a, 24b, and the memory controllers 26a, 26b (Fig. 2) of the CPUs 12A, 12B as single units, since that is the way they function.

Reintegration is used to place two CPUs in... both of the paired CPUs at virtually the same time.

The major steps in the process for changing from simplex mode operation of the one on-line CPU to duplex mode... greater detail by the flow diagrams of Figs. 33A - 33D, generally are:

1. Setup and synchronize the two CPUs (one on-line, the other off-line) and their connected routers to... the memory of the on-line CPU to the off-line CPU, maintaining a tracking process that monitors changes in the memory of the on-line CPU that have not been... and may need to be copied over to, the off-line CPU;
3. Setup and synchronize the CPUs to run a delayed (slave) duplex mode from the same instruction stream (lock... will write the predetermined registers (not shown) of the control registers 74 in the interface units 24 of CPUs 12A and 12B, to a next state (after a soft operation) in... the off-line CPU 12B.

Next, a sequence is entered (steps 1060 - 1070) that will synchronize the clock synchronization FIFOs of the CPUs 12A, 12B and routers 14A, 14B in much the same fashion... the same steps described above in connection with the discussion of Figs. 31A, 31B to synchronize the clock synchronization FIFOs. The on-line CPU 12A will send the sequence of a SLEEP symbol, self-addressed message packet, and SYNC symbol which, with the SYNC CLK signal, operates to synchronize CPUs and routers. Once so synchronized, the on-line CPU 12A then, at step 1066, sends a Soft Reset (SRST) command... of all configuration registers and control registers (e.g., configuration registers 74 of the interface units 24) cache, and the like to memory 28 of the on-line CPU, ...time to have the system 10 off-line for reintegration. For that reason, the reintegration process is performed in a manner that allows the on-line CPU to continue executing user... not match that of the off-line CPU. The reason for this is that normal processing by the processor 20 of the on-line CPU can change memory content after it has been copied... when a memory location is written in the on-line CPU 12A during the reintegration process it is marked as "dirty;" second,

all copying of memory to the off-line CPU... ..may, however, limit the ability to detect two-bit errors. But, since the memory copying process will last for a only relatively short period of time, this risk is believed acceptable... ..memory location in CPU 12A is made (either an incoming I/O write, or a processor write operation). The returning data (that was copied over to the off-line CPU) would... ..controller 26 (Fig. 2) of the on-line CPU to monitor memory locations in the process of being copied over to the off-line CPU 12B. The memory controller uses a...within the block had been written by another operation (e.g., a write by the processor 20, an I/O write, etc.), that prior write operation will flag the location in... ..still must be copied over to the off-line CPU 12B.

Returning to the reintegration process, and now to Fig. 33B, the memory tracking (AtomicWrite mechanism and using ECC to mark... ..entails writing a reintegration register (not shown; one of the configuration registers 74 of interface unit 24 - Fig. 5) to cause a reintegration (REINT) signal to be asserted. The REINT signal... ..is left alone. Throughout the incremental copy operations, the normal actions of the on-line processor will mark some memory locations dirty.

Several passes of incremental copying will need to be... ..the number of successful WriteConditional operations at the end of each pass through memory, the processors 20 can determine the effect of a given pass compared to the previous pass. When the benefits drop off, the processors 20 will give up on the precopy operations. At this point the reintegration process is ready to place the two CPUs 12A, 12B into lock-step operation.

Thus, the... ..in Fig. 33C, where at step 1100, the on-line CPU 12A momentarily halts foreground processing, i.e., execution of a user application. The remaining state (e.g., configuration registers, cache, etc.) of the on-line processors 20 and its caches is then read and written to a buffer (series of memory... ..to the off-line CPU 12B, together with a "reset vector" that will direct the processor units 20 of both CPUs 12A, 12B to a reset instruction.

Next, step 1106 will quiesce... ..to ensure that the FIFOs of the routers are clear, that the FIFOs of the processor interfaces 24 are clear, and no further incoming I/O message packets are forthcoming. At... ..symbol will be received and acted upon by both CPUs 12A, 12B, to cause the processor units 20 of each CPU to jump to the location in memory 28 containing the reset... ..a subroutine that will restore the stored state of both CPUs 12A, 12B to the processor units 20, caches 22, registers, etc. The CPUs 12A, 12B will then begin executing the same... ..enabling of the ECC bit to mark dirty locations must now be disabled, since the processors are doing the same thing to the same memory. During this stage of the reintegration... ..encountered by CPU 12A.

Meanwhile, the bus error in the CPU 12A will cause the processor unit 20 to be forced into an error-handling routine to determine (1) the cause of... ..error was caused by an attempt to read a memory location marked dirty. Accordingly, the processor unit 20 will initiate (via the BTE 88 -- Fig. 5) the AtomicWrite mechanism to copy the...the SRST symbols are now received by the CPUs 12A, 12B, they will cause both processor units 20 of the CPUs to be reset to start from the same location with the... ..will periodically update, e.g., a database or audit file that is indicative of the processing of the primary CPU up to that point in time of the update. Should the... ..in error-checking redundancy to the CPU 12B, in the same manner that the individual processor units 20a, 20b of the CPU 12A provide fail-fast, fault tolerance for the CPU - when... ..cost system is applicable, as illustrated in Fig. 34. As shown in Fig. 34, a processing system 10' includes the CPU 12A and routers 14A, 14B structured as described above. The... ..and the CPUs are also the same.

Thus, the CPU 12B' comprises only a single processor unit 20' and associated support components, including the cache 22', interface unit (IU) 24', memory controller 26', and memory 28'. Thus, while the CPU 12A is structured in the manner shown in Fig. 2, with cache processor unit, interface unit, and memory control redundancies, approximately one-half of those components are needed to implement CPU... ..stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... ..i.e., parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of...
...inputs connected to receive the data output

from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... ...for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the...
...attached, without the assistance of any other CPU 12 in the system. Many prior parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O device becomes... ...rest of the system. Other prior systems provide access to I/O through pairs of processors so that should one of the processors fail, access to the ...if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... ...multiprocessing system access to every peripheral, as done here, operates to extend the "primary"/"backup" process taught in the above-identified U.S. Patent No. 4,228,496. There, a multiple CPU system may have a primary process may running on one CPU, while a backup process resides in the background on another of the CPUs. Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be... ...e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system... ...CPU can access any I/O without the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... ...four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... ...an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time...indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... ...When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... ...response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Specification: ...stream. CPU 12A is designed to provide fail-fast operation through the duplication of the processor unit 20 and other elements that make up the CPU. In addition, through the duplex operation... ...i.e., parity checks at various interfaces), data integrity is missing.

Fig. 34 illustrates the processing system 10' as including a pair of routers 14A, 14B to perform the comparing of...
...inputs connected to receive the data output

from the CPUs 12A and 12B' have clock synchronization FIFOs as described above to receive the somewhat asynchronous receipt of the data output, pulling... ...for the moment to Figs. 1A-1C, an important feature of the architecture of the processing system illustrated in these Figures is that each CPU 12 has available to it the...

...attached, without the assistance of any other CPU 12 in the system. Many prior parallel processing systems provide access to or the services of I/O devices only with the assistance of a specific processor or CPU. In such a case, should the processor responsible for the services of an I/O device fail, the I/O device becomes... ..rest of the system. Other prior systems provide access to I/O through pairs of processors so that should one of the processors fail, access to the corresponding I/O is still available through the remaining I/O... ..if both fail, again the I/O is lost.

Also, requiring the resources of a processor in order to provide any other processor of a parallel or multi-processing system imposes a performance impact upon the system.

The ability to allow every CPU of... ..multiprocessing system access to every peripheral, as done here, operates to extend the "primary"/"backup" process taught in the above-identified U.S. Patent No. 4,228,496. There, a multiple CPU system may have a primary process running on one CPU, while a backup process resides in the background on another of the CPUs. Periodically, the primary process will perform a "check-pointing" operation in which data concerning the operation of the process is stored at a location accessible to the backup process. If the CPU running the primary process fails, that failure is detected by the remaining CPUs, including the one on which the backup resides. That detection of CPU failure will cause the backup process to be activated, and to access the check-point data, allowing the backup to resume the operation of the former primary process from the point of the last check-point operation. The backup process now becomes the primary process, and from the pool of CPUs remaining, one is chosen to have a backup process of the new primary process. Accordingly, the system is quickly restored to a state in which another failure can be... ..e., failed CPU) has been repaired.

Thus, it can be seen that the method and apparatus for interconnecting the various elements of a the processing system 10 provides every CPU with access to every I/O element of that system the necessity of using the services of another processor. Thereby, system performance is enhanced and improved over systems that do require a specific processor to be involved in accessing I/O.

Further, should a CPU 12 fail, or be... ..four bit Transaction Sequence Number (TSN) field; see Figs. 3A and 3B. Elements of the processing system 10 (Fig. 1) which are capable of managing more than one outstanding request, such... ..an expected response to a prior issued request message packet bound for an I/O unit 17 or a CPU 12 is not received within a predetermined allotted period of time... ..indicate a fault in the communication path. An interrupt will be generated internally, and the processors 20 (20a, 20b - Fig. 2) will initiate execution of a barrier request (BR) routine. That... ..When the Barrier Request message packet (i.e., 1150) is received by the X interface unit 16a of the I/O packet interface 16A, it will formulate a response message packet... ..response to the barrier request message packet is received by the CPU 12A it is processed through the AVT logic 90' (see also Figs. 5 and 11). The barrier response uses...

Claims: ...A2

1. In a computing system including first and second processor elements of substantially identical construction coupled to one another for communicating data therebetween, each of the first and second processor elements including a memory element for storing instructions and data, the first processor element operating to execute instructions of an instruction stream from the memory element, a method for establishing for synchronized, substantially lock-step operation of the second processor element with the first processor element to have the second processor element executing the same instructions at substantially the same moment in time as the first processor element, the method including the steps of:

synchronizing the second processor element with the first processor element;

the first processor element accessing the instructions and data of the memory element of the first processor element and communicating the accessed instructions and data to the second processor element with address data indicative of locations in the memory element of the second processor element corresponding to locations at which the instructions and data are stored in the memory element of the first processor element;

storing the received instructions and data in the memory element of the second processor element at the locations indicated by the address; and

periodically sending selected ones of the instructions and data of the first processor

element for storing in the memory element of the second processor element at locations corresponding to location of the first memory element.

2. The method of claim 1, wherein the synchronizing step includes delaying operation of the second processor element relative to the operation of the processor element.

3. The method of claim 1, including the step of marking first locations of the memory element of the first processor that are written with new data or new instructions after the instructions and data accessed from the first locations for communication to the second processor element.

4. The method of claim 3, wherein the data and instructions each include a... ..of marking.

6. The method of claim 1, wherein each of the first and second processor elements is coupled to a pair of router elements for at least interprocessor communication, the synchronizing step including setting the router elements to accept communication only from the first processor element.

7. The method of claim 2, wherein each of the first and second processor elements is coupled to a pair of router elements for at least interprocessor communication, the synchronizing step includes delaying communication of data from the pair of router elements to the second processor element relative to communication of data from the pair of router elements to the first processor element.

Claims: ...B1

1. Procède pour établir un fonctionnement synchronisé, sensiblement en lock-step, d'un premier processeur (12A) avec un deuxième processeur (12B) pour que le deuxième processeur exécute les mêmes instructions sensiblement au même moment que le premier processeur, les premier et deuxième processeurs ayant une construction sensiblement identique et étant couplés l'un... ..instructions et des données, le procédé comprenant les étapes consistant à :

synchroniser (1050) le deuxième processeur (12B) avec le premier processeur (12A) ;

mettre en oeuvre le premier processeur pour exécuter les instructions d'un flot d'instructions provenant de la première mémoire, le premier processeur accédant aux instructions et aux données de la première mémoire et communiquant les instructions et les données auxquelles il accède au deuxième processeur avec des données d'adresse indicatives des emplacements dans la deuxième mémoire correspondant aux emplacements... ..maintenir, dans une mémoire, une table des emplacements de la première mémoire communiqués au deuxième processeur ;

mémoriser (1084) les instructions et les données reçues dans la deuxième mémoire aux emplacements indiqués... ..et/ou données à partir des premiers emplacements en vue de leur communication au deuxième processeur ;

confirmer une communication réussie d'instructions et de données à la deuxième mémoire en utilisant... ..1086) des instructions et des données sélectionnées parmi les instructions et les données du premier processeur en vue de leur mémorisation dans la deuxième mémoire à des emplacements correspondant à des... ..dans lequel l'étape de synchronisation comprend le fait de retarder le fonctionnement du deuxième processeur par rapport au fonctionnement du premier processeur de sorte que le deuxième processeur fonctionne un certain nombre de périodes d'horloge après le premier processeur pour exécuter une instruction identique à une instruction exécutée par le premier processeur le certain nombre de périodes d'horloge auparavant.

3. Procède selon la revendication 1, dans... ...synchronisation comprenant le paramétrage des routeurs pour accepter une communication uniquement en provenance du premier processeur.

6. Procède selon la revendication 2, dans lequel chacun des premier et deuxième processeurs est... ...comprenant le fait de retarder la communication de données des deux routeurs vers le deuxième processeur par rapport à la communication de données des deux routeurs vers le premier processeur.

7. Procède selon la revendication 1, dans lequel, pendant l'étape de copie des emplacements chaque emplacement de la mémoire du premier processeur qui n'est pas marqué comme étant << propre >>.

15. Procède selon la revendication 14, comprenant... ...étant << sales >> en écrivant de nouvelles données et/ou instructions auxdits emplacements par le premier processeur.

16. Procède selon la revendication 15, comprenant en outre l'étape consistant à arrêter temporairement...

6/K/30 (Item 30 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Country	Number	Kind	Date		
Type		Pub. Date		Kind	Text
Available Text		Language		Update	Word Count
Total Word Count (Document A)					
Total Word Count (Document B)					
Total Word Count (All Documents)					

Specification: ...is set ON again. 9.2.5. Fault in Timer

The DS3-SMDS interface performs processes such as the monitor of the performance of the DS3-PLCP layer based on the... ...trigger is entered within 15 minutes + 15 seconds after the preceding input timing, then static processes such as the performance monitor process, etc. cannot be performed. Therefore, if a trigger is not entered on a predetermined schedule... ...again. Since the fault point is accommodated according to the hardware monitor, no special software process is required.

9.2.6. DS3 Layer Alarm

The DS3-SMDS interface monitors the carrier... ...interface and the switch software. The switch software refers to the program executed by the processor which controls the processes (call process, switch control process, etc.) of the entire switch.

10.2. Hardware Interface

As shown in Figures 8 and... ...switch path by way of the MDX and ASSWSH. The BSGCSH communicates with the switch processor through an interface (INF).

The extraction/insertion of an intra-station control communications cell from... ...communications data transmitted through the ASSWSH of the standby system is discarded by the common unit of the BSGCSH in the standby

system. The common unit of the BSGCSH in the standby system identifies an intra-office control communications cell by... ..path (intra-switch highway). In this case, the DS3-SMDS interface performs a dropping/inserting process only on an intra-station communications cell input/output upwards (at ASSWSH). No dropping/inserting processes are performed on an intra-station communications cell input/output via the line (DS3 transmission... ..the DS3-SMDS interface loaded to the RMXSH in the BRLC performs a dropping/inserting process only on an intra-station communications cell input/output upwards (at the station), and no dropping/inserting processes are performed on an intra-station communications cell input/output via the subscriber line. Therefore, the DS3-SMDS interface passes an intra-station communications cell transferred from a downward unit to the BSGCSH.

The intra-station communications cell between the DS3-SMDS interface and the... ..the BSGC to the DS3-SMDS interface is added by the VCC in the common unit of the BSGC.

10.4 Error Monitor

The DS3-SMDS interface does not directly monitor... ..its tag as a valid intra-station communications cell addressed to the interface, and then processes the cell.

10.5. AAL Interface

10.5.1. SAR-PDU Format

Figure 67 shows... ..refer to the 4.2.2. and 4.2.3. in part 3). The AAL process performed by the DS3-SMDS interface has the functions of (1) decomposing/composing an L2... ..If a bit error is detected in the payload of a cell in the AAL process, then the cell is discarded. The error is stored in the DS3-SMDS interface and... ..MSCN. If an SN error or an ST sequence error is detected in the AAL process, then a series of cells determined to be erroneous are all discarded. In the AAL process, accepted as valid cells are those related to the SSM without payload errors, or a... ..SMDS interface and displayed as an MSCN. No detected errors are corrected in the AAL process.

10.8. L2 Interface

10.8.1. Functions of L2

A simple LAP is the... ..3) Setting the downward DMUX-LSI buffer threshold (when necessary)

11.2. Blocking

The following processes are performed.

(1) Setting Block Specification (OUS)

11.3. Setting In-Service

The following processes are performed.

(1) Resetting the block specification (OUS)

(2) Setting/resetting master reset (M-RST... ..the E-MSCN service

(5) Transferring various initialization data

11.4. Non-implementation

The following processes are performed.

(1) Setting Block Specification (OUS)

11.5. Processes for faults

11.5.1. Monitor of Faults

A fault of the DS3-SMDS interface... ..points on the specific condition to each point.

11.5.2. Detection of faults

The processes to be performed when each of the representative NG-OR points is detected are listed... ..by referring to another area of the MSCN or by directly inquiring of the individual unit through the intra-station control communications.

(1) At the detection of a hardware fault;

1... ..cells are discarded in buffer) Since the MSCN displays data based on a predetermined statistics process in the hardware, messages are displayed based on the displayed data.

11.5.3. Specifying a fault

(1) When the ASSWSH is processed as OUS;

A fault is specifies by automatic diagnostics of a faulty ASSWSH system.

(2... ..an ASSWSH system is switched and the diagnostics is manually carried out. A series of processes are manually performed. The online diagnostics refers to the diagnostics actually performed by a switch processor (CC) of an active system regardless of the state of the DS3-SMDS interface.

11.5.4. Monitor of Recovery

(1) ASSWSH and DS3-SMDS interface

These units are recovered when they are changed from the OUS state to the INS state. If... ..monitored for faults.

(2) Line System Alarm

An MSCN monitor constantly monitors the recovery of units. If no blocking factors exist at the time of recovery, a blocked DS3-SMDS interface... ..in Buffer Recovery is constantly monitored through the monitor of the MSCN.

11.6 Various Process Sequence

Figures 70 through 81 show the sequence of the following processes.

(1) Initialization of DS3-SMDS interface

(2) Procedure of INS of DS3-SMDS interface

(3... ..intra-station control communications

2. Hardware fault which disables intra-station control communications

3. Micro processor fault

4. Communications error between the SIFSH common and the DS3-SMDS interface (active system... ..the SIFSH common and the DS3-SMDS interface (standby system)

(5) DS3/PLCP layer alarm process

(6) Notification of D/Q Timer (counting every 15 minutes and every day) at the...self-diagnostics function, the normality of the hardware of the simplex portion (excluding the communications unit) of the DS3-SMDS interface can be confirmed.

Listed below are the steps of the... ..interface.

(1) Initialization

(2) Checking the SRAM

(3) Checking the dual port RAM (simple LAPD process)

(4) Read/write check of each LSI loaded on the DS3-SMDS interface

(5) Pseudo... ..detection and notification system for each fault state as being associated with each fault correction process in the DS3-SMDS interface loaded in the subscriber interface shelf (SIFSH).

14.1.1... ..Fault of disconnected fuse

(4) Fault of erroneous insertion of package

(5) Fault of individual unit package (fault of simplex unit)

14.1.2 OBP Fault

In the SIFSH, power-through packages are loaded separately on... ..power is supplied for a half shelf independently.

14.1.3. OBP Fault in Individual Unit (DS3-SMDS interface)

A fault of an OBP (power source) loaded on the DS3-SMDS interface 1 is detected in the SIFSH common (SIF-COM, common unit) in both active and standby systems. The fault is detected by monitoring the display of the individual unit OBP fault register in the SIFSH common and the occurrence of a stack in the E-MSCN highway.

An output of the LED output terminal unit of the OBP indicates an open state in a normal operation and a ground state in an abnormal operation. When an output of the LED terminal unit indicates the ground state, a fault value is set in the OBP fault register.

Figure 88 shows the configuration of the OBP monitoring function in the individual unit.

(1) +5V OBP Fault

If a +5V OBP fault has arisen in the DS3-SMDS interface individual unit, then a serial highway for the extended maintenance scanner (E-MSCN) information to be provided... ..is blocked with a stack. There are representative points indicating the IDs of the individual units in the E-MSCN, and the occurrence of a stack for the points is monitored... ..and standby systems. The fault is actually detected by monitoring the display of the individual unit OBP fault register in the SIFSH common and the occurrence of a stack in the E-MSCN highway. Each individual unit comprises a plurality of packages. If there is a package missing among a plurality of... ..the +5V power source to be provided in the entire package group in the individual unit is not induced. Accordingly, the SIFSH common monitors the items indicating the ID point of the individual unit in the E-MSCN toward the SIFSH common to detect all "H" (high level) for... ..of the systems, then it determines that an interface fault has occurred between the individual unit and the SIFSH common. The state is checked when the systems are switched.

Figure 89... ..configuration of the package missing monitoring function.

14.1.5. Fuse Disconnection Fault

The individual unit fuse provided for the power package is individually monitored in the SIFSH common of both... ..missing fault to be detected because a highway stack simultaneously occurs in a corresponding individual unit. However, a fuse disconnection fault is detected by priority by the firmware in the SIFSH... ..Fault

In the SIFSH, a package group comprising a plurality of packages in the individual unit and the SIFSH common can have the configuration in which the OBP can be activated... ..packages and their circuit elements are not destroyed.

14.1.7. DS3-SMDS Interface Individual Unit Package Fault

There are following two types of hardware faults of a package in the DS3-SMDS interface individual unit.

(1) Hardware fault notified of through the intra-office control communications using the E-MSCN... through 63 as being related to the faults defined by (1) above.

1. MPE (micro-processor fault)

2. FEER-1 (fault indicating that the intra-station control communications cannot be established... fault)

4. UHDPT (upward highway data parity error fault)

5. EGPTY (intra-station control communications terminal LSI fault) Next, listed below are the points in the C-MSCN shown in Figures... the L2-PDU cell has 8-bit width of parallel data. The DS3-SMDS interface

processes 16-bit width of parallel data at the transmission speed of 9.72 Mbps. Therefore... written to the dual port RAM. However, the MPE bit in the 017th byte is processed by the hardware.

Data are sequentially read from the dual port RAM using as an... 15.1.1.8. Microprocessor Interface Function

The HAFOOA PCB is loaded with the 80C186 processor and outputs processor interface signals of the HAFOOA and other PCBs.

15.1.2. Functions of HLPO1A

The most important function of the HLPP1A (Figure 45) is to perform a process specific to the DS3-SMDS.

Among the DS3-SMDS interface functions described in 7., the... data conversion function

(2) 45Mbps --> 156Mbps data conversion function

(3) Distributed queue dual bus (DQDB) process function The outline of these functions is explained below. Figure 92 shows the configuration of... data conversion function is realized by the V2 DMUX LSI.

15.1.2.3. DQDB Process Function

This function is explained in 7.6. above.

14.1.3. Functions of HDTOOA... Firmware Interface

16.1. General Descriptions

The DS3-SMDS interface is loaded with the 80C186 processor to realize the following functions.

(1) DS3 layer performance monitor

(2) PLCP layer performance monitor... the DS3-SMDS interface is realized using the control chip select (CS) from the 80C186 processor.

The control chip select conditions in each interface are listed below, and Figure 93 is... The subscriber interface shelf type A (SIFSH-A) can be loaded with up to 8 units per shelf of the individual units containing the ATM subscriber interface circuits.

The following 5 types of the individual units can be accommodated.

(1) OC3C (156Mbps optical interface unit) (simplex configuration)

(2) DS-3 (45Mbps metallic interface unit) (simplex configuration)

(DS3-SMDS interface explained in Part 2)

(3) ADSINF (ADS1SH concentrator unit) (duplex configuration)

(4) TCGADP (TCGSH adapter unit) (simplex configuration: two systems of the TCGSH are connected to a single unit)

(5) LOOP (156Mbps loop unit) (duplex configuration) Each unit of the OC3C, DS-3, and TCGADP has a simplex configuration. Each unit of the ADSINF and LOOP has a duplex configuration. If the units are mounted to the SIFSH-A, then a two-unit set is accommodated. Accordingly, up to 4 sets of the ADSINF and LOOP units can be loaded per shelf.

The active/standby control for each unit of the ADSINF and LOOP can be performed by the SIFSH common unit (hereinafter referred to as the SIFCOM).

If the SIFSH-A (SIFSH) is mounted to the... 94, then the SIFSH-A functions as shelf exclusive to the load of the LOOP unit. If the SIFSH-A is mounted to the left of the ASSW (ATM switch) in Figure 94, then the SIFSH-A functions as shelf for loading the individual unit which terminates a subscriber line.

The SIFCOM in the SIFSH-A performs the intra-station signalling process to the broadband signaling group controller shelf (BSGC) connected to the ASSW through the BSGCSH. The BSGC converts the command issued by the switch software and executed by the switch processor (CC) (not shown in the drawings) by way of the interface type T (INFT) into... through the INFT.

A simple LAP-D protocol is adopted to the intra-station signalling process. The simple LAP-D protocol is developed to minimizing the function of the hardware and firmware based on the LAP-D protocol.

Among the individual units accommodated in the SIFSH-A, each unit of the OC-3C and DS-3 communicates with the BSGC using the simple LAP... command in an EMSD highway if the analysis result indicates a command to an individual unit, and notifies the individual unit of the result.

The SCN information from the individual unit is multiplexed in time divisions in the EMSCN highway and notified to the SIFCOM. The... whose data change is detected.

The SIFCOM demultiplexes an ATM cell corresponding to each individual unit from the downward cell highway which has a transmission speed of 622Mbps and is connected... cell highway which has a transmission speed of 156Mbps and is connected to each individual unit.

The ATM cell in the 156Mbps upward cell highway connected to each individual unit is multiplexes in the 622Mbps cell highway connected to the ASSW. A scheduler system is... later in 6.1.2. The scheduler system multiplexes an upward cell from each individual unit in the arrival order such that the order can be maintained correctly in both active... switched in the ASSW and SIFCOM.

The SIFSH-A can accommodate up to 8 individual units per shelf. However, to improve the multiplexing of cells from the 156Mbps highway to the... INS is incorporated

Passing/conversion variable mode of ATM cell having 0 bit

(4) Individual unit interface

Transmitting and receiving cells in a 156Mbps cell highway

Generating and checking the parity... in a 156 Mbps cell highway

Passing/discard control of a cell from an individual unit of a standby system (monitoring 0 bit)

Detecting an individual unit missing

Specifying the slot number of an individual unit

Specifying active/standby switching for a duplex device (MUXACTD signal)

Notifying of completion of active...by-cell/collective selection available)

Preventing a corresponding test cell from flowing to an individual unit at the loopback of a test cell

Various self-diagnostics

(11) Power source

-48V 5 system/one-way supply

Loading each SIFCOM and individual unit with an onboard power module (OBP)

Automatic power down of the SIFCOM of a corresponding... ...A is 3 (steps).

2.1. Configuration

Described below are the SIFCOM and each individual unit.

2.1.1. SIFCOM

The SIFCOM is fixedly loaded on the SIFSH-A and is... ...per system as shown in Figure 95. The HPT01A package in the SIFCOM provides each unit in a single system with a -48V power source. Each of the systems on the... ...left of the center of the shelf is power-supplied separately.

2.1.2. Individual Unit

Up to 8 individual units can be loaded on the SIFSH-A.

Each individual unit is composed of 3 packages per unit. The names of slots accommodating these packages are slots A, B, and C from left... ...shelf. -48V/CG is power-supplied independently from the power through package to each individual unit and SIFCOM. The power through package is loaded with a maintainer fuse corresponding to each individual unit and SIFCOM. The CG is independently connected to each of the systems on the right... ...risk plate.

2.2.3. +5V/E

+5V is provided in each of the individual units. The earth E is shared among systems 0 and 1.

The power sources -48V/CG... ...Interface

Described below are the interface and signal timing between the SIFSH-A and other units.

3.1. Switch Interface

The SIFSH-A comprises a 622Mbps cell highway and an interface... ...state, an alarm state, and a selected system state in each system.

3.3. Individual Unit Interface

Described below are the interface and signal timing between the SIFCOM and individual unit loaded on the SIFSH-A through the back-wiring board (BWB). All interface points between the SIFCOM and individual units explained below are defined according to the polarity and timing in the BWB.

3.3.1. 156Mbps cell highway interface

The interface of the 156Mbps cell highways between the common unit and the individual unit is explained below.

As shown in Figure 102, the ATM cell in the 156 Mbps... ...the timing of receiving an ATM cell from the upward cell highway from the individual unit to the SIFCOM. The individual unit transmits an upward cell by receiving a

cell request signal from the SIFCOM because the... ..through the scheduler at the SIFCOM requires the upward cells from each circuit to be synchronized.

3.3.1.2. Downward 156Mbps Cell Highway Interface

Figure 104 shows the timing of receiving an ATM cell from the downward cell highway from the SIFCOM to the individual unit. The SIFCOM transmits a downward cell by receiving a cell request signal from the individual unit so that the downward cell frame can be synchronized in the SIFCOM of both systems to prevent the generation of duplicate or missing cells in fetching a downward cell in each individual unit in a downward cell fetching process.

3.3.2. E-MSD/E-MSCN Highway Interface

The physical and logical specifications are described below for the EMSD/EMSCN highway between the SIFCOM and individual unit.

The downward (SIFCOM --> individual unit) data highway is defined as an EMSD highway. The EMSD is transferred to the SIFCOM... ..the simple LAP-D, multiplexed in the EMSD highway, and serially transferred to the individual unit.

The upward (individual unit --> common unit) data highway is defined as an EMSCN highway. The EMSCN is an echo-back (EMSD normally received at the individual unit and looped back to the EMSCN highway) to the EMSD, and fault status information in the individual unit. The EMSCN is multiplexed in the EMSCN highway and serially transferred to the SIFCOM. A... ..LAP-D communications.

3.3.2.1. System Control

An internal circuit in the individual unit operates according to the EMSD, CLK, and FCK from the SIFCOM of an active system... ..circuit of an ACT controller. The circuit which receives an ACT0/ACT1 in the individual unit is necessarily pulled up so that an "L" active control can be performed in both... ..MSB to LSB; the downward EMSD data highway and the upward EMSCN data highway are synchronized in bit and byte position

Hard reset (HRST) : individual unit hard reset signal; reset with "1" in the BWB and output asynchronously

Fault reset (FRST): individual unit fault reset signal; reset with "1" in the BWB and output asynchronously

3.3.2.3. Logical Specification

3.3.2.3.1. Individual Unit Receiving Specification

Described below is the logical specification of the EMSD receiving process in the individual unit.

The receiving terminal is protected against SIFCOM interface fault (noises of the EMSD, etc., stack fault, etc.) by frame synchronization, checking a pilot signal, and twice reading

processes.

Figure 112 is a flowchart showing the operations of these processes. Figure 113 is a block diagram showing the functions of the individual unit for performing these processes in series.

3.3.2.3.2. Frame Synchronization

The frame synchronization corresponds to step 1 shown in Figure 112 and the functional portion 1 shown in Figure 113.

The number of protection steps for the frame synchronization of the EMSD highway is 1 step each for forward and backward. The stack of... ..both L/H stacks) are detected.

Figure 109 shows the state transition of the frame synchronization process.

Practically, data is fetched from a corresponding frame when a normal synchronization FCK is received in a hunting state as shown in Figure 110. If an abnormal FCK is once received in a synchronization established state, then the frame synchronization state changes into the hunting state and the data are discarded from this point, but the data received immediately before the point is stored until the synchronization is established next time. A normal FCK refers to the fact that the receiving terminal counter value (for example, a carry-out) depending on the CLK/FCK matches the next... EMSD is the 000th byte/bit D7 (refer to Figures 58 and 59).

The individual unit detects an EMSD highway stack fault when the alternation of the pilot signals 0/1 becomes irregular. The individual unit discards the data at and after an abnormal point as shown in Figure 111, and... D6 (PLTF); refer to Figures 61 and 62.

3.3.2.3.4. Twice Reading Process

The data fetched in the frame synchronization process described in the 3.3.2.3.2. and the pilot 0/1 signal check process described in 3.3.2.3.3. is stored in a noise erase memory 4... in Figure 112). If these data do not match, then they are discarded.

A protection process is performed using a DTEN signal (step 4 shown in Figure 112). The DTEN signal... indicate "L" in the BWB by a microprocessor in the SIFCOM. When the intra-shelf units are turned on simultaneously, a rise time conflict occurs after the release of the power-on reset for the SIFCOM and the individual unit, and a value of the EMSD highway becomes uncertain. The DTEN signal is used to control the individual unit such that it cannot fetch the EMSD data. Therefore, the individual unit ignores all EMSD data when the DTEN signal indicates "H". The DTEN signal is accommodated... the EMSD highway (refer to Figures 58 and 59).

3.3.2.3.5. Individual Unit Sending Specification

Described below is the logical specification of an EMSCN sending process in the individual unit.

The EMSCN of an active system transmits an echo-back in response to the EMSD... information.

Figure 114 is a block diagram showing the EMSCN sending circuit in the individual unit.

3.3.2.3.6 Fault Detection

Figure 115 is a list of the methods of detecting and notifying in the individual unit of the interface fault between the SIFCOM and the individual unit, and of the method of detecting the fault in the SIFCOM and the contents of... 4. Clock Interface

The clock interface refers to clock systems in the SIFCOM and individual unit along the flow of cells.

In the SIFCOM, a cell is written to the DMUX... obtained by dividing a 77.76MHz clock transferred from the ASSW (ATM switch) into 6 units.

As shown in Figure 116, a cell is read from the DMUX buffer in the DMX-LSI to the individual unit in synchronism with a 19MHz (19.44MHz precisely) clock transferred from the individual unit. The 19MHz clock from the individual unit is generated as follows. That is, as shown in Figure 116, a 64KHz clock is transferred to the individual unit in the SIFCOM after being obtained by dividing into 128 units an 8MHz clock received from the SYNSH through an optical link. According to the clock, the PLL module in the individual unit generates a 156MHz (155.52MHz precisely) clock. Then, the above described 19MHz clock can be... also generates a 156MHz clock according to the 64KHz clock obtained by dividing into 128 units the 8MHz clock received from the SYNSH. An upward cell is written to the MUX... LSI corresponding to each circuit in synchronism with the 19MHz clock transferred from the individual unit. The cell is read from the MUX buffer in synchronism with the 13MHz

(12.96MHz... 10.9 of part 2. The switch software refers to a program executed in a processor for controlling the entire process of the switch (call process, switch control process, etc.).

4.1. Outline

The SIFCOM communicates with the switch software by performing an intra-switch passing through the ASSWSH (refer to Figure 94). The BSGC communicates with the switch processor through an interface type T (INFT).

A simple LAP-D is a protocol newly developed... reduce the load on the hardware and firmware. Specifically, numbered frames in layer 2, which charge a heavy load on the hardware, can be successfully removed. As a result, only unnumbered frames are processed in layer 2. To avoid missing and duplicate messages, numbered frames are processed in layer 3. Since the number management function is originally indispensable for firmware, the numbered... BOM) for its ST and 44 bytes for its LI. The SAR-PDU storing an intermediate segment is assigned a continuation of message (COM) for its ST and 44 bytes for... 2. above.

5. Allocation of Tag

Figure 121 shows the format of an ATM cell processed in the SIFSH-A.

According to the present embodiment, an ATM cell is routed using... ASSW the ATM cell (whose header has been converted by a VCC) transferred from individual units #0 (sup(equivalent to)) #7 accommodated in the SIFSH-A, and a signalling cell generated by the signal processing unit in the SIFCOM.

If the SIFSH is connected in series, then the multiplexing control of... MUX.

The MUX multiplexes a cell in the 156Mbps upward highway connected to each individual unit and a signalling cell generated in the signal processing unit (shown in Figure 130) in the SIFCOM in the 622Mbps upward highway to the ASSW. The cell transferred from each individual unit is input to the MUX after its header is converted according to the VCC (refer to Figure 130).

The MUX comprises a buffer for 52 cells corresponding to each individual unit, and only valid cells are stored in the buffer. Each buffer notifies the multiplexing control unit (scheduler) of a write of a cell each time a cell is written to the... The multiplexing control of an ATM cell in the 156Mbps highway extended from each individual unit is performed by a scheduler. A scheduler is assigned to each 622Mbps upward highway. If... the 156Mbps highway has been written to the buffer is transmitted from a write control unit (not shown in Figure 133) in each buffer to the scheduler.

As shown in Figure... scheduler contains a FIFO having 18-bit width corresponding to the number of circuits (individual units) to be monitored by the scheduler, samples the write completion signal received from each circuit... corresponds to the time required to transmit one cell in the 600Mbps highway.

Each individual unit has a simplex configuration, while the SIFCOM has a duplex configuration. This scheduler multiplexing control... 52 cells (8 bits x 54 octet x 52 cells = 22464 bits) per circuit (individual unit) as a buffer used in multiplexing ATM cells in a low-speed input highway into... Congestion control is not performed (refer to 6.1.9.).

6.1.6. Abnormal Write Process

If an abnormal cell described in 6.1.6.1. and 6.1.6.2. below is input, then the following abnormal write process is performed.

6.1.6.1. Too small cell length

If the data length of... inputs "H" indicating an output permission signal to the buffer.

6.1.8. Abnormal Read Process

If the scheduler inputs to a buffer an output permission signal at intervals within approximately... ..in series into a cell toward a low-speed highway downward each of the individual units in the SIFSH-A and a signaling cell input to the signal processing unit in the SIFCOM. The cells are demultiplexed according to the tag in the header of... ..in the DMUX.

The DMUX demultiplexes a cell to each of up to 8 individual units in the shelf and a signaling cell from the 622Mbps high-speed highway according to... ..the DMUX transmits the former through the 156Mbps low-speed highway connected to each individual unit, and the latter to the signal processing unit (Figure 140) in the SIFCOM. In this case, the DMUX comprises a buffer for 112 cells for each individual unit.

A cell dropper (cell DRP) for each individual unit in the DMUX shown in Figure 141 determines whether or not a cell is dropped... ..in the SIFCOM according to command A at the DMUX 0 corresponding to the individual unit accommodating the active circuit of the umbilical link as shown in Figure 144, while TAGC = "000" is set according to command B at the DMUX 4 corresponding to the individual unit accommodating a standby circuit of the umbilical link. If a fault occurs in the active...the maximum buffer length, which is an initial value, is set as a cell discard process start threshold.

Listed below are the relationships between each threshold and a buffering operation in... ..the SIFCOM on the following grounds.

Assume that the VCC is loaded into the individual unit having the configuration of a duplex VCC. Furthermore, assume that the cell transmitted from subscriber... ..the assumption above, further assume that a fault occurs at the VCC in the individual unit for subscriber A (A sub) as shown in Figure 146, and that a cell is... ..subscriber line may undesirably affect 64 or more circuits.

In this case, the fault detecting process can monitor an MC (monitoring cell) at a receiving equipment. In this process, a fault can be detected by inserting a monitoring cell (MC1 and MC2 in Figure... ..ASSW are switched. However, since the fault has occurred in the VCC of the individual unit having a simplex configuration, a switch fault will soon occur in a newly active ASSW... ..the future virtual path (VP) services.

Table 1 (Table-1) is used to retrieve an intermediate VPI using an input VPI (VPI assigned to an input cell) as an address. According to the present embodiment, an input VPI value = an intermediate VPI value assuming that no VP services are provided.

Table 2 (Table-2) is used to retrieve an output VPI/VCI using an intermediate VPI + input VCI (VCI assigned to an input cell) as an address. According to the present embodiment, an input VPI value = an intermediate VPI value assuming that no VP services are provided.

6.3.3. Inter-System VCC... ..table and is generated at a writing operation.

6.3.3.4. Procedure for INS process

The state transition from OUS to INS is carried out after a switch processor (CC) issues a copy start command to instruct the VCC table of an active system... ..CC issues a reset request command to the SIFCOM of the OUS system. The copy process is performed after the contents of the VCC table in the SIFCOM of the OUS... ..the CC of the reset completion notification status after the reset is completed. The reset process enables only the VPI/VCI on the VCC table in the SIFCOM of the active... ..the copying time.

Figure 148 is an arrow diagram showing the procedure for an INS process. The procedure is described below by referring to Figure 148.

If the copy process terminates normally, then the SIFCOMs of both systems notify the CC of a copy completion status. Unless the copy process terminates normally due to an inter-system communications fault, etc. from no response of a... ..is provided for the CC. As a result, the CC determines failure in the copy process and resets again the SIFCOM of the OUS system. If any of the SIFCOMs of... ..OUS system is reset again. Figure 149 shows the status of each system and the process of the CC.

Normally, a set/release command (call process command) is issued by the CC to the SIFCOM of both systems independently. The SIFCOM is configured such that it can receive a call process command in a VCC copy process. During the VCC copy process, the command is issued by the CC not to the SIFCOM of both systems but to the SIFCOM of the active system. This is because the call process command reaches the SIFCOM of the OUS system faster than the SIFCOM of the active... ..of the OUS system may be set again to the previous contents through the copy process on the VCC table from the SIFCOM of the active system when the VCC table... ..through the hardware complicates the protocol and enlarges the scale of the hardware, a call process command is issued only to the SIFCOM of the active system.

Accordingly, if the state... ..the operation state, then required is a protocol for preventing the specification of a call process command from the CC to the SIFCOM of the old OUS system from being lost... ..system.

(3) The SIFCOM of the active system copies to the other system all call process commands received before receiving the command described in (2) above. All call process commands received after receiving the command described in (2) above are executed only in its... ..command to the SIFCOM of the OUS system.

(7) If the queue stores a call process command to a new standby system while the processes described in (3) through (6) are executed, then the CC issues the command immediately.

After the process (7) above, the CC issues a call process command independently to each SIFCOM of the active and standby systems.

6.3.3.5... ..service from the specific value (for example, VPI=3F, VCI=03FF) added by the individual unit, and simultaneously changes the value of the VPI/VCI assigned to the header of the ATM cell containing the payload field input from the individual unit of the DS3-SMDS interface, etc. and the L2-PDU of the SMDS service into the value of the VPI/VCI specifying the subscriber network interface (SNI) terminating the individual unit which transmits the ATM cell. Accordingly, the PVC established between the SIFCOM and the SBMESH... ..VCI of the number corresponding to the number of the SNI terminated by the individual unit connected to the SIFCOM and used for the SMDS service. The SIFCOM adds to the... ..ATM cell in the ATM switch and transferring it to the SBMESH.

6.4. Signaling Process (EGCLAD)

6.4.1. Outline

Figure 150 shows the position of the signal processing unit (EGCLAD) in the SIFSH-A.

An EGCLAD LSI converts between a simple LAP-D-based... ..shown in Figure 153, an MC is inserted in a subscriber interface at an input terminal. The MC should be inserted at predetermined intervals of cells for each path. The SINF at an output terminal requires the function of monitoring the MC inserted at predetermined cell intervals.

Monitoring an MC... ..through the ASSW of a standby system are discarded in the SIFCOM at the output terminal of the standby system and do not reach the SINF at the output terminal as indicated by broken lines shown in Figure 153.

Accordingly, the quality of a path... ..and SIFCOM are loaded with the cell-by-cell loopback function for performing a normal process on a user cell and looping back only cells generated by the TCG.

The cell... ..function indicates a loopback for each VPI/VCI. Therefore, the switch software notifies a loopback unit of a VPI/VCI value of a looped-back cell through an MSD.

Since the... ..function which is activated by the MSD information from the switch software.

8. Fault Correcting Process

8.1. Fault Detection Point and Notification System

Described below is the fault detection and... ..fault

(4) SIFCOM package front connector missing fault

(5) Package erroneous insertion fault

(6) Individual unit package fault (simplex unit fault)

(7) SIFCOM package fault (duplex unit fault)

a) Individual unit interface fault

b) Common unit fault

(8) Individual unit-SIFCOM interface fault (simplex/duplex cross-connected portion fault)

8.1.2. OBP Fault

This fault is described in 14.1.2. in part 2.

8.1.2.1. Individual Unit OBP Fault

This fault is described in 14.1.3. in part 2.

8.1... ..fault monitor object system as shown in Figure 155.

The output of the LED output terminal of the OBP indicates a release state in a normal operation and a ground state... ..fault value is set in the OBP fault register when the output of the LED terminal indicates a ground state.

Since the SIFCOM comprises 4 packages and each package is loaded with an OBP, a signal line connecting the LED output terminals of all these OBP is connected to the SIFCOM of the mate system.

8.1.3. Package Missing Fault

8.1.3.1. Individual Unit Package Missing Fault

This fault is described in 14.1.4. in part 2.

8... ..shown in Figure 157.

8.1.4. Fuse Disconnection Fault

8.1.4.1. Individual Unit Fuse Disconnection Fault

This fault is described in 14.1.5. in part 2.

8... ..the switch software.

(2) Lower Order Shelf --> ASSW

As shown in Figure 160, the detecting unit similar to that shown in Figure 159 relating to (1) above is mounted to both... ..This fault is described in 14.1.6. in part 2.

8.1.7. Individual Unit Package Fault

This fault is described in 14.1.7. in part 2.

8.1... ..Fault

The faults in the SIFCOM are classified into the following two types.

(1) Interface unit fault in individual unit

(2) Common unit fault Figure 162 shows the component in which a fault occurs. Figure 163 shows a... ..broadband remote line concentrator (BRCL; refer to Figure 34) or a broadband remote line switching unit (BRSU) and a host switch.

According to the present embodiment, two in-band signaling routes... ..2. Line Reassignment Sequence

All faults in an umbilical line are detected in the individual unit (OC3C or DS-3; refer to Figure 94).

A detected line fault is provided as EMSCN information by the individual unit for the SIFCOM, and then transmitted from the SIFCOM to the switch software via the... ..information is read according to a command request from the switch software to the individual unit.

The individual unit notifies the switch software of the detailed fault information in response to the command.

Figure 165 shows the sequence of reassigning a line in a line protection process.

9.3. Setting VCC in Standby Line

A standby line is provided with a VCC... ..As shown in Figure 166, this command only contains the information about the identification number (Unit No.) of a unit which changes a tag value and about the tag value (TAGC) itself. That is, a... ..of control data between the SWMDX (HMX03A), SWMX (HSR00A), or SCLK (HTG02A) and the switch processor (CC).

3. Interface

3.1. Communication Line System

Figure 168 shows the connection configuration of... ..excluding an enable signal. Parity bits of valid cells only are checked in the input unit of the ATM switch, and parity bits are assigned to valid cells only in the output unit of the ATM switch. The contents of the data in the information field (payload) of... ..and standby systems. Each block in the SWCNT and ASSWSH-A is connected via a processor data bus and an address bus.

Each block is controlled mainly by monitoring a fault... ..control system of the ASSWSH-A can be an active/standby control function to each terminal unit. As shown in Figures 167 and 168, the SWCNT comprises 32 output units corresponding to 32 output highways of 622Mbps on both ends (sides 0 and 1: left and right sides of the SWMX) through the SWMDX. From the output units through the SWTIF not shown in Figure 167, a system selection signal and its strobe... ..system, it is output as a signal having the same polarity in both systems. Each terminal unit selects an active system device in the system according to the system selection logic shown...from the non-assured services. Therefore, only the P bit is used to control the process. During the congestion, a cell designated for a non-assured service is discarded.

5.2... ..Figure 192 is assigned in the ASSWSH-A to the 2.4Gbps/622 Mbps DMUX unit in the SWMX and the SWMDX. A threshold (Xp) is set for the cell buffer... ..of the SWMX.

5.2.2. Congestion Control in SWMDX

The 2.4Gbps/622Mbps DMUX unit in the SWMDX is provided in the ADMUX LSI shown in Figure 182. Setting the... ..time the occurrence of the cell discard is reported to the CC and the reporting processes are different between the SEMX and SWMDX. The cell discard reporting processes for the SWMX and SWMDX are described individually as follows.

In the SWMX, cell discard... ..SWMDX, cell discard is not regarded as a fault. Since the 622Mbps/2.4Gbps MUX unit in the SWMDX is an STM, the no discard occurs and the discard portion is exclusively the 2.4Gbps/622Mbps DMUX unit. The number of times in every 15 minutes that cells are discarded is counted in the traffic measure process described in 5.3. The occurrence of cell discard is recognized by the CC's reading the count value.

5.3. Traffic Measure Process

In the ASSWSH-A, the number of the following cells is counted in the 2.4Gbps/622Mbps DMUX unit as the function similar to the performance monitor in order to manage the status of... ..the 4 lower bits in the 4th word at the address. Described below is the process in the ASSWSH-A at the reception of each order.

- Activation of a command: The... ..SWCNT and response in a specified format in a data bus.

6.3. Fault Correcting Process

6.3.1. Fault Detection

The important functions of the firmware in the SWCNT are... ..the fault in the MSCN, then the MSCN table is also updated.

The above described process is realized by the process modules (1) through (3) listed below.

(1) Alarm interruption handler

Trigger of process occurrence of a fault reading a fault register updating a fault counter generating fault notification data (MSG BOX)

activating a fault correcting task

(2) Cycle activation task

Trigger of process 100msec cycle

Process comparing fault counters clearing fault counters

(3) Fault correcting task

Trigger of process receiving an MSG BOX

Process notifying a higher order process of the contents of a fault:

generating detailed fault data updating an MSCN table generating... ..6.3.2. Message Box

Figure 202 shows a basic format of a message box processed by the fault correcting task.

(1) Listed below are the contents of the message box... ..4. Self-diagnosis

Upon receipt of a self-diagnosis setup command from the CC (switch processor), the firmware in the SWCNT makes a diagnosis of each fault monitoring function according to... ..the following orders among the orders described in 6.2. above and performs a diagnostic process and checks the result.

(1) SWMX compulsory alarm highway parity error

(2) SWMX compulsory alarm... ..it is not mounted for the slot which returned no answer.

The firmware performs these processes only for load-recognized slots.

(2) After terminating the initialization of the device, the state... ..into the operating state in which the system initialization is performed from a higher order process. At this time, the firmware is notified of the load state of the HMX03A displayed... ..and the state is compared with the load state recognized by the firmware in the process described in (1) above.

(3) In the comparing process in (2) above, if there is a slot which is recognized by the firmware as... ..the MSCN and includes the slot in the detailed fault data.

(4) In the comparing process described in (2) above, if there is a slot which is recognized by the firmware... ..case, the subsequent control is performed according to the station information.

7.3. Fault Correcting Process

The ASSWSH-A has the specification of monitoring a fault as follows.

(1) A duplex... ..is adopted as a redundant configuration (one shelf for one system)

(2) Various fault detection processes are performed, and the systems are switched according to the detection result (control by the... ..shelf (SBMESH) switches data of the SMDS subscriber. The switch is performed actually in cell unit while the message format is checked. As for a protocol, terminated are level 2 (AAL... ..subscriber message handler (SBMH) as shown in Figure 204.

In Figure 204, an actual SMDS terminal unit is connected beyond the subscriber network interface (SNI). Likewise, an switching system (SS) is connected... ..and an R portion, and the data input from the SNI to the system is processed in the S portion of the SBMESH, and the data processed in the R portion of the SBMESH (SBMH) is output from the system to the... ..the GWMESH (WGMH) is described in Part 6.

1.1.2. Outline of SMDS Data Process

Figure 205 shows the route of SMDS data between SNIs, and the data is processed in the following procedure.

1. The data input from the SNI to the ASSW (UP... ..represented by bold lines.

Thus, in the case of the data transmission between the SNIs, processes are performed only by the SBMH. When the data is transferred to and from other SS and LATA SS, the processes are performed by the SBMH and GWMH. The actual routing control, the relationship between each... ..diagram showing the SBMESH.

As shown in Figure 209, the SBMESH comprises an MH-COM unit for interfacing with the ASSW and an LP unit for performing actual switching.

The MH-COM unit comprises an SDMX, RDMX, SMUX, and RMUX. The characters S and R for the MUX... ..not show in Figure 209. The VCC is set by the LAP. The MH-COM unit has a checking function and detected information is provided with interface to the software through the LAP or the broadband signaling controller (BSGC) described later in Part 7.

The LP unit comprises an SMLP, RMLP, and LP-COM. The initial characters S and R of the... ..required for a switching, subscriber data, information detected by each checking function in the LP unit, billing information, etc. are provided with interface to the software through the INF.

As described above... ..and demultiplexed by the SDMX, RDMX, SMUX, and RMUX. On the other hand, the LP unit and the INF is connected one to one. For example, if four SBMESHs are daisy... ..accordingly.

1.3. Redundant Configuration

As shown in Figure 210, the MH-COM and LP units have duplex configurations (systems #0 and #1).

The MH-COM unit has a master/slave configuration exclusive for the ASSW, while the LP unit has an independent master/slave configuration. The master system (for example, #0) and slave system (for example, #1) of the LP unit have basically the same function, and the slave system can actually perform a switching operation. In this case, the billing information obtained through the slave system's switching is not reported to the software.

There is an inter-system cross-connection between the duplex MH-COM unit and LP unit, that is, between system #0 of the MH-COM unit and system #1 of the LP unit and between system #1 of the MH-COM unit and system #0 of the LP unit. However, no inter-system cross-connection exists between system #0 of the LP unit and system #1 of the INF and between system #1 of the LP unit and system #0 of the INF.

The RMLP in system #0 of the LP unit receives data from the RDMX of system #0 of the MH-COM unit and data from the RDMX of system #1 of the MH-COM unit. The selector (not shown in the figure) in the input unit of the RMLP selects the data from the master system of the MH-COM unit. Likewise, the SMUX of the MH-COM unit receives data from the SMLP of system #0 of the LP unit and data from the SMLP of system #1 of the LP unit. The selector (not...

6/K/31 (Item 31 from file: 348)

EUROPEAN PATENTS

(c) 2008 European Patent Office. All rights reserved.

Systeme et methode pour la communication entre des processus

Country	Number	Kind	Date
---------	--------	------	------

Abstract ...information from the data portion of the message allows the performance of the message passing process to be linked to the relative complexity of the message to be transferred between two... subsystem to the destination task. If the message cannot be transmitted, for example because of processor resource exhaustion, a time out expiration, or insufficient port rights, then processor time is not wasted in the abortive transfer of the data portion of the message...

Type	Pub. Date	Kind	Text
Available Text	Language	Update	Word Count
Total Word Count (Document A)			
Total Word Count (Document B)			
Total Word Count (All Documents)			

Specification: ...A2

The present invention relates to improvements in operating systems for data processing systems.

The invention disclosed herein is related to the copending United States Patent Application by... Sotomayor, Jr., James M. Magee, and Freeman L. Rawson, III, which is entitled "METHOD AND APPARATUS FOR MANAGEMENT OF MAPPED AND UNMAPPED REGIONS OF MEMORY IN A MICROKERNEL DATA PROCESSING SYSTEM", Serial Number 263,710, filed June 21, 1994, IBM Docket Number BC9-94-053... Patent Application by James M. Magee, et al. which is entitled "CAPABILITY ENGINE METHOD AND APPARATUS FOR A MICROKERNEL DATA PROCESSING SYSTEM", Serial Number 263,313, filed June 21, 1994, IBM Docket Number BC9-94-071... Patent Application by James M. Magee, et al. which is entitled

"TEMPORARY DATA METHOD AND APPARATUS FOR A MICROKERNEL DATA PROCESSING SYSTEM", Serial Number 263,633, filed June 21, 1994, IBM Docket Number BC9-94-076... ..by James M. Magee, et al. which is entitled "MESSAGE CONTROL STRUCTURE REGISTRATION METHOD AND APPARATUS FOR A MICROKERNEL DATA PROCESSING SYSTEM", Serial Number 263,703, filed June 21, 1994, IBM Docket Number BC9-94-077... ..Application by James M. Magee, et al. which is entitled "ANONYMOUS REPLY PORT METHOD AND APPARATUS FOR A MICROKERNEL DATA PROCESSING SYSTEM", Serial Number 263,709, filed June 21, 1994, IBM Docket Number BC9-94-080... ..Srinivasan, Dennis Ackerman, and Himanshu Desai which is entitled "PAGE TABLE ENTRY MANAGEMENT METHOD AND APPARATUS FOR A MICROKERNEL DATA PROCESSING SYSTEM", Serial Number 303,805, filed September 9, 1994, IBM Docket Number BC9-94-073... ..Gupta, Ravi Srinivasan, Dennis Ackerman, and Himanshu Desai which is entitled "EXCEPTION HANDLING METHOD AND APPARATUS FOR A MICROKERNEL DATA PROCESSING SYSTEM", Serial Number 303,796, filed September 9, 1994, IBM Docket Number BC9-94-072... ..Rawson, Ching-Yun Chao, and Charles Jung, which is entitled "BACKING STORE MANAGEMENT METHOD AND APPARATUS FOR A MICROKERNEL DATA PROCESSING SYSTEM", Serial Number 303,851, filed September 9, 1994, IBM Docket Number BC9-94-087... ..copending United States Patent Application by Ching-Yun Chao, et al., which is entitled "MASTER SERVER PROGRAM LOADING METHOD AND APPARATUS FOR A MICROKERNEL DATA PROCESSING SYSTEM", Serial Number 308,189, filed September 19, 1994, IBM Docket Number BC9-94-074... ..incorporated herein by reference.

The operating system is the most important software running on a computer. Every general purpose computer must have an operating system to run other programs. Operating systems typically perform basic tasks... ..not access the system.

Operating systems can be classified as multi-user operating systems, multi-processor operating systems, multi-tasking operating systems, and real-time operating systems. A multi-user operating... ..same time. Some operating systems permit hundreds or even thousands of concurrent users. A multi-processing program allows a single user to run two or more programs at the same time. Each program being executed is called a process. Most multi-processing systems support more than one user. A multi-tasking system allows a single process to run more than one task. In common terminology, the terms multi-tasking and multi-processing are often used interchangeably even though they have slightly different meanings. Multi-tasking is the... ..at the same time, a task being a program. In multi-tasking, only one central processing unit is involved, but it switches from one program to another so quickly that it gives...systems use preemptive multi-tasking, whereas the Multi-Finder (TM) operating system for Macintosh (TM) computers uses cooperative multi-tasking. Multi-processing refers to a computer system's ability to support more than one process or program at the same time. Multiprocessing operating systems enable several programs to run concurrently. Multi-processing systems are much more complicated than single-process systems because the operating system must allocate resources to competing processes in a reasonable manner. A real-time operating system responds to input within a short... ..determines to a great extent the applications which can be run. For IBM compatible personal computers, example operating systems are DOS, OS/2 (TM), AIX (TM), and XENIX (TM).

A user... ..commands are accepted and executed by a part of the operating system called the command processor or command line interpreter.

There are many different operating systems for personal computers such as CP/M (TM), DOS, OS/2 (TM), UNIX (TM), XENIX (TM), and AIX (TM). CP/M was one of the first operating systems for small computers. CP/M was initially used on a wide variety of personal computers, but it was eventually overshadowed by DOS. DOS runs on all IBM compatible personal computers and is a single user, single tasking operating system. OS/2, a successor to DOS, is a relatively powerful operating system that runs on IBM compatible personal computers that use the Intel 80286 or later microprocessor. OS/2 is generally compatible with DOS... ..supports virtual memory. UNIX and UNIX-based AIX run on a wide variety of personal computers and work stations. UNIX and AIX have become standard operating systems for work stations and are powerful multi-user, multi-processing operating systems.

In 1981 when the IBM personal computer was introduced in the United States, the DOS operating system occupied approximately 10 kilobytes of storage. Since that time, personal computers have become much more complex and require much larger operating systems. Today, for example, the OS/2 operating system for the IBM personal computers can occupy as much as 22 megabytes of storage. Personal computers become ever more complex and powerful as time goes by and it is apparent that... ..The goal of that research was to develop a new operating system that would allow computer programmers to exploit modern hardware architectures emerging and yet reduce the size and the number... ..thread, the port, the message, and the memory object.

The task is the traditional UNIX process which is divided into two separate components in the MACH microkernel. The first component is... ..ports. A task is a passive collection of resources; it does not run on a processor.

The thread is the second component of the UNIX process, and is the active execution environment. Each task may support one or more concurrently executing...set of services for building operating system personalities implemented as a set of user-level servers. The Microkernel System 115 is made up of many server components that provide the various traditional operating system functions and that are manifested as operating system personalities. The Microkernel System 115 uses a client/server system structure in which tasks (clients) access services by making requests of other tasks (servers) through messages sent over a communication channel. Since the microkernel 120 provides very few services... ..how to manage the interprocess communication that must take place between the many clients and servers in the system, in a fast and efficient manner.

In accordance with the present invention... ..interprocess communication in a microkernel architecture, the system comprising: a memory means in a data processing system, for storing data and programmed instructions; a data bus means coupled to said memory means in said data processing system, for transferring signals; a processor means coupled to said memory means with said data bus means, for executing said programmed... ..buffer for storing message data information, and having a first thread executing instructions in said processor means, for forming a first message to send to a destination port; a second task... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means; and a transmission control means in said interprocess communications means, for interpreting said message... ..aspect, there is now provided a system for interprocess communications in a microkernel architecture data processing system, comprising: a memory in the data processing system, for storing information; a program for a first task in said memory, that includes... ..template for the first task, that includes a pointer to the transmission control buffer; a processor means associated with a thread of the first task, for executing the instructions from the program; said processor means executing a first instruction in the thread, to load a data value into the send data buffer and to load a control value into the transmission control buffer; said processor means executing the procedure call in the thread, to make the header template available to... ..interprocess communication in a microkernel architectures, comprising: storing in a memory means in a data processing system, data and programmed instructions; executing in a processor means coupled to said memory means, said programmed instructions; coordinating in an interprocess communications means... ..buffer for storing message data information, and having a first thread executing instructions in said processor means, for forming a first message to send to a destination port; storing a second... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means, and interpreting with a transmission control means in said interprocess communications means, said message...communication with another task in a microkernel architecture, comprising: a memory means in a data processing system, for storing data and programmed instructions; an application program means in said memory means, for providing application program instructions to be executed; a processor means coupled to said memory means, for executing said programmed instructions; a microkernel means in... ..buffer for storing message data information, and having a first thread executing instruction in said

processor means, for forming a first message to send to a destination port; a second task... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means, and a transmission

control means in said interprocess communications means, for interpreting said message... ..communication with another task in a microkernel architecture, comprising: a memory means in a data processing system, for storing data and programmed instructions; an operating system personality program means in said memory means, for providing operating system personality program instructions to be executed; a processor means coupled to said memory means, for executing said programmed instructions; a microkernel means in... ..buffer for storing message data information, and having a first thread executing instructions in said processor means, for forming a first message to send to a destination port; a second task... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means; and a transmission control means in said interprocess communications means, for interpreting said message... ..communication with another task in a microkernel architecture, comprising: a memory means in a data processing system, for storing data and programmed instructions; a personality-neutral services program means in said memory means, for providing personality-neutral services program instruction to be executed; a processor means coupled to said memory means, for executing said programmed instructions; an interprocess communications means... ..buffer for storing message data information, and having a first thread executing instructions in said processor means, for forming a first message to send to a destination port; a second task... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means; and a transmission control means in said interprocess communications means, for interpreting said message... ..provided a system 'as claimed in any of claims 23 to 25', comprising: a second processor means coupled to said memory means, for executing said programmed instructions; a third thread in... ..associated with said second task, for providing said programmed instructions for execution in said second processor means.

Viewing the present invention from a tenth aspect, there is now provided a system... ..communication with another task in a microkernel architecture, comprising: a memory means in a data processing system, for storing data and programmed instructions; an operating system personality program means in said memory means, for providing operating system personality program instructions to be executed; a processor means coupled to said memory means, for executing said programmed instructions; a microkernel means in... ..buffer for storing message data information, and having a first thread executing instructions in said processor means, for forming a first message to send to a destination port; a second task... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means; and a transmission control means in said interprocess communications means, for interpreting said message... ..communication with another task in a microkernel architecture, comprising: a memory means in a data processing system, for storing data ...in said memory means, for providing personality-neutral services program instructions to be executed; a processor means coupled to said memory means, for executing said programmed instructions; a microkernel means in... ..buffer for storing message data information, and having a first thread executing instructions in said processor means, for forming a first message to sent to a destination port; a second task... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means, and a transmission control means in said interprocess communications means, for interpreting said message... ..aspect, there is now provided a method for interprocess communications in a microkernel architecture data processing system, the method comprising the steps of: loading a program for a first task into... ..memory, associated with the first task, for executing the instructions from the program in a processor; executing a first instruction in the thread with a processor, to load a data value into the send data buffer and to load a control value into the transmission control buffer; executing the procedure call in the thread with the processor, to make the header template available to an interprocess communications subsystem; and establishing the transmission... ..aspect, there is now provided a system for interprocess communications in a microkernel architecture data processing system, comprising: a memory in the data processing system, for storing information; a program for a first task in said memory, that includes... ..template for the first task, that includes a pointer to the transmission control buffer; a processor means associated with a thread of the first task, for executing the instructions from the program; said processor means executing a first instruction in the thread, to load a data value into the send data buffer and to load a control value into the transmission control buffer; said processor means executing the procedure call in the thread, to make the header template available to... ..aspect, there is now provided

a system for interprocess communications in a microkernel architecture data processing system, comprising: a memory means in the data processing system, for storing information; a program for a first task in said memory means, that... ..transmission control buffer; a data bus means coupled to said memory means in the data processing system, for transferring signals; a processor means coupled to said memory means with said data bus means and associated with a thread of said first task, for executing the instructions from said program; said processor means executing a first instruction in the thread, to load a data value into said... ..control buffer; an interprocess communications subsystem in said memory means, for managing message transfers; said processor means executing the procedure call with the thread, to make said header template available to... ..In one aspect the present invention thus provides an improved microkernel architecture for a data processing system. The improved microkernel architecture of the present system is more simple in its interprocess... ..faster and more efficient interprocess communication capability. Still furthermore, the microkernel architecture for a data processing system, that has greater flexibility in the exchange of messages between tasks within a shared memory environment and between distributed data processors that do not share a common memory.

The above advantages and other objects, features and... ..a preferred embodiment of the present invention by the separation of transmission control method and apparatus for a microkernel data processing system disclosed herein.

In even a moderately complex multitasking application, many tasks and threads are...message, in a preferred embodiment of the invention, allows the performance of the message passing process to be linked to the relative complexity of the message to be transferred between two... ..port of the destination task. If the message cannot be transmitted, for example because of processor resource exhaustion, a time out expiration, or insufficient port rights, then processor time is not wasted in the abortive transfer of the data portion of the message...and yet it maximizes the performance of the system.

The client task and/or the server task can be part of an application program, an operating system personality program, a personality... ..communicate with still other tasks concurrently running on different host multiprocessor systems in a distributed processing network. Each communication from one such task to another can avail itself of the efficiencies... ..the invention manages the interprocess communication that must take place between the many clients and servers in a Microkernel System, in a fast and efficient manner. The invention applies to uniprocessors, shared memory multiprocessors, and multiple computers in a distributed processor system.

Preferred embodiments of the present invention will now be described by way of example... ..8 shows a functional block diagram of two host multiprocessor systems running in a distributed processing arrangement, with the IPC subsystem and the transmission control module on each host processor managing interprocess communications between tasks with the exchange of messages between the two hosts over...support through the capability engine.

Figure 28 illustrates the basic execution loop of the multiplexing server.

Figure 29 is the message passing library anonymous reply algorithm.

Figure 30 illustrates share region... ..run multiple operating system personalities 150 on a variety of hardware platforms.

The host multi-processor 100 shown in Figure 1 includes memory 102 connected by means of a bus 104... ..devices, or other I/O devices. Also connected to the bus 104 is a first processor A, 110 and a second processor B, 112. The example shown in Figure 1 is of a symmetrical multi-processor configuration wherein the two uni-processors 110 and 112 share a common memory address space 102. Other configurations of single or multiple processors can be shown as equally suitable examples. The processors can be, for example, an Intel 386 (TM) CPU, Intel 486 (TM) CPU, a Pentium (TM) processor, a Power PC (TM) processor, or other uni-processor devices.

The memory 102 includes the microkernel system 115 stored therein, which comprises the microkernel 120, the personality neutral services (PNS) 140, and the personality

servers 150. The microkernel system 115 serves as the operating system for the application programs 180... of the machine. The microkernel system 115 includes the microkernel 120 and a set of servers and device drivers that provide personality neutral services 140. As the name implies, the personality neutral servers and device drivers are not dependent on any personality such as UNIX or OS/2. They depend on the microkernel 120 and upon each other. The personality servers 150 use the message passing services of the microkernel 120 to communicate with the personality neutral services 140. For example, UNIX, OS/2 or any other personality server can send a message to a personality neutral disc driver and ask it to read... to be built by adding pieces to the smaller ones. For example, each personality neutral server 140 is logically separate and can be configured in a variety of ways. Each server runs as an application program and can be debugged using application debuggers. Each server runs in a separate task and errors in the server are confined to that task.

Figure 1 shows the microkernel 120 including the interprocess communications... IPC) 122, the virtual memory module 124, tasks and threads module 126, the host and processor sets 128, I/O support and interrupts 130, and machine dependent code 125.

The personality... 140 shown in Figure 1 includes the multiple personality support 142 which includes the master server, initialization, and naming. It also includes the default pager 144. It also includes the device... support and device drivers. It also includes other personality neutral products 148, including a file server, network services, database engines and security.

The personality servers 150 are for example the dominant personality 152 which can be, for example, a UNIX personality. It includes a dominant personality server 154 which would be a UNIX server, and other dominant personality services 155 which would support the UNIX dominant personality. An alternate ...be for example OS/2. Included in the alternate personality 156 are the alternate personality server 158 which would characterize the OS/2 personality, and other alternate personality services for OS... the Microkernel System 115 carefully splits its implementation into code that is completely portable from processor type to processor type and code that is dependent on the type of processor in the particular machine on which it is executing. It also segregates the code that... drivers; however, the device driver code, while device dependent, is not necessarily dependent on the processor architecture. Using multiple threads per task, it provides an application environment that permits the use of multi-processors without requiring that any particular machine be a multi-processor. On uni-processors, different threads run at different times. All of the support needed for multiple processors is concentrated into the small and simple microkernel 120.

This section provides an overview of... the following features:

Support for multiple personalities

Extensible memory management

Interprocess communication

Multi-threading

Multi-processing

The Microkernel System 115 provides a concise set of kernel services implemented as a pure... set of services for building operating system personalities implemented as a set of user-level servers.

Objectives of the Microkernel System 115 include the following:

Permit multiple operating system personalities to... extensible communication kernel;

An object basis with communication channels as object references; and

A client/server programming model, using synchronous and asynchronous inter-process communication.

The basis for the Microkernel System 115 is to provide a simple, extensible communication... ..Support of address spaces for tasks; and

Management of physical resources, such as physical memory, processors, interrupts, DMA channels, and clocks.

User mode tasks implement the policies regarding resource usage. The... ..a C runtime environment, including such basic constructs as string functions, and a set of servers which include:

Name Server - Allows a client to find a server

Master Server - Allows programs to be loaded and started

Kernel Abstractions

One goal of the Microkernel System... ..This can make it difficult to identify key ideas. The main kernel abstractions are:

Task - Unit of resource allocation, large access space and port right

Thread - Unit of CPU utilization, lightweight (low overhead)

Port - A communication channel, accessible only through the send/receive capabilities or rights

Message - A collection of data objects

Memory object - The internal unit of memory management (Refer to Section 2, Architectural Model, for a detailed description of the... ..concepts).

Tasks and Threads

The Microkernel System 115 does not provide the traditional concept of process because: All operating system environments have considerable semantics associated with a process (such as user ID, signal state, and so on). It is not the purpose of the microkernel to understand or provide these extended semantics.

Many systems equate a process with an execution point of control. Some systems do not.

The microkernel 120 supports multiple points of control separately from the operating system environment's process.

The microkernel provides the following two concepts:

Task

Thread

(Refer to Section 2, Architectural Model...against system paging space.

Task to Task Communication

The Microkernel System 115 uses a client/server system structure in which tasks (clients) access services by making requests of other tasks (servers) through messages sent over a communication channel. Since the microkernel 120 provides very few services... ..execute in a virtual environment. The virtual environment provided by the kernel contains a virtual processor that executes all of the user space accessible hardware instructions, augmented by user-space PNS and emulated instructions (system traps) provided by the kernel. The virtual processor accesses a set of virtualized registers and some virtual memory that otherwise responds as does...libraries consist of stubs that invoke the microkernel's IPC system to send messages to servers. This architecture permits the flexible implementation of function: servers can be replaced by other servers and services can be combined into single tasks without affecting

the sources of the programs... ..to its clients and to elements of the PNS. Thus, the dominant personality is a server of "last resort". The dominant personality implements whatever services are defined by the PNS libraries but are not implemented by another server.

The microkernel 120 is also dependent on some elements of the PNS. There are cases when it sends messages to personality-neutral servers to complete internal kernel operations. For example, in resolving a page fault, the microkernel 120... ..a virtual environment. The virtual environment provided by the microkernel 120 consists of a virtual processor that executes all of the user space accessible hardware instructions, augmented by emulated instructions (system traps) provided by the kernel; the virtual processor accesses a set of virtualized registers and some virtual memory that otherwise responds as does... ..and a set of threads.

Security Token:

A security feature passed from the task to server, which performs access validations.

Port:

A unidirectional communication channel between tasks.

Port Set:

A set of ports which can be treated as a single unit when receiving a message.

Port Right:

Allows specific rights to access a port.

Port Name... ...It is through this object that the memory manager manipulates the clients' visible memory image.

Processor:

A physical processor capable of executing threads.

Processor Set:

A set of processors, each of which can be used to execute the threads assigned to the processor set.

Host:

The multiprocessor as a whole.

Clock:

A representation of the passage of time...resources. The kernel provides communication methods that allow a client task to request that a server task (actually, a thread executing within it) provide a service. In this way, a task... ..underscore)trap. This trap allows the thread to send messages to the kernel and other servers to operate upon resources. This trap is almost never directly called; it is invoked through... ..token is included as an implicit value in all messages sent by the task. Trusted servers can use this sent token as an indication of the sender's identity for use in making access mediation decisions.

A task inherits the security token of its parent. Because this token is to requests a service and a server that provides the service. A port has a single receiver and potentially multiple senders. The... ..A port is a unidirectional communication channel between a client who requests service and a server who provides the service. If a reply is to be provided to such a service... ..an object basis. Some operations require two objects, such as binding a thread to a processor set. These operations show the objects separated by commas. Not all entities are named by... ..A port set is a set of ports that can be treated as a single unit when receiving a message. A mach(underscore)msg receive operation is allowed against a port...ports. The operations upon abstract memory objects include the following:

Initialization

Page reads

Page writes

Synchronization with force and flush operations

Requests for permission to access pages

Page copies

Termination

Memory...the memory manager have been disposed Restrict access to memory pages

Provide performance hints

Terminate

Processor

Each physical processor that is capable of executing threads is named by a processor control port. Although significant in that they perform the real work, processors are not very significant in the microkernel, other than as members of a processor set. It is a processor

set that forms the basis for the pool of

processors used to schedule a set of threads, and that has scheduling attributes associated with it. The operations supported for processors include the following:

Assignment to a processor set

Machine control, such as start and stop

Processor Set

Processors are grouped into processor sets. A processor set forms a pool of processors used to schedule the threads assigned to that processor set. A processor set exists as a basis to uniformly control the schedulability of a set of threads. The concept also provides a way to perform coarse allocation of processors to given activities in the system. Processor sets are characterized by a uniformity with respect to the scheduling of threads which run... ..on one or another without emulation or migration of its environment. The operations supported upon processor sets include the following:

Creation and deletion

Assignment of processors

Assignment of threads and tasks

Scheduling control

Host

Each machine (uniprocessor or multiprocessor) in a... ..following:

Clock manipulation

Statistics gathering

Re-boot

Setting the default memory manager

Obtaining lists of processors and processor sets

Clock

A clock provides a representation of the passage of time by incrementing a... thread. Also, faults or other illegal instruction behaviour cause the kernel to invoke its exception processing.

Figure 2. shows the client visible structure associated with a thread. The thread object is... thread port is also accessible as the thread's thread self port, through the containing processor set or the containing task.

Reference is made here to the above cited copending United... Guy G. Sotomayor, Jr., James M. Magee, and Freeman L. Rawson, III, entitled "METHOD AND APPARATUS FOR MANAGEMENT OF MAPPED AND UNMAPPED REGIONS OF MEMORY IN A MICROKERNEL DATA PROCESSING SYSTEM", which is incorporated herein by reference for its... be derived from the task's task self port, the contained threads or the containing processor set.

Reference is made here to the above cited copending United States Patent Application by Guy G. Sotomayor, Jr., James M. Magee, and Freeman L. Rawson, III, entitled "METHOD AND APPARATUS FOR MANAGEMENT OF MAPPED AND UNMAPPED REGIONS OF MEMORY IN A MICROKERNEL DATA PROCESSING SYSTEM", which is incorporated herein by reference for its more detailed discussion of this topic... port is a unidirectional communication channel between a client who requests a service and a server who provides the service. A port has a single receiver task and can have multiple... Guy G. Sotomayor, Jr., James M. Magee, and Freeman L. Rawson, III, entitled "METHOD AND APPARATUS FOR MANAGEMENT OF MAPPED AND UNMAPPED REGIONS OF MEMORY IN A MICROKERNEL DATA PROCESSING SYSTEM", which is incorporated herein by reference for its more detailed discussion of this topic... Guy G. Sotomayor, Jr., James M. Magee, and Freeman L. Rawson, III, entitled "METHOD AND APPARATUS FOR MANAGEMENT OF MAPPED AND UNMAPPED REGIONS OF MEMORY IN A MICROKERNEL DATA PROCESSING SYSTEM", which is incorporated herein by reference for its more detailed discussion of this topic... virtual memory system is designed for uniform memory access multiprocessors of a moderate number of processors. Support for architectures providing non-uniform memory access or no remote memory access is currently... Guy G. Sotomayor, Jr., James M. Magee, and Freeman L. Rawson, III, entitled "METHOD AND APPARATUS FOR MANAGEMENT OF MAPPED AND UNMAPPED REGIONS OF MEMORY IN A MICROKERNEL DATA PROCESSING SYSTEM", which is incorporated herein by reference for its more detailed discussion of this topic... Guy G. Sotomayor, Jr., James M. Magee, and Freeman L. Rawson, III, entitled "METHOD AND APPARATUS FOR MANAGEMENT OF MAPPED AND UNMAPPED REGIONS OF MEMORY IN A MICROKERNEL DATA PROCESSING SYSTEM", which is incorporated herein by reference for its more detailed discussion of this topic... manages interprocess communications between three tasks 210, 210", and 210' with threads running on two processors 110 and 112. The data processing system can be a shared memory, multiprocessing system as is shown in Figure 7, a distributed processing system as is shown in Figure 8, or a uniprocessor system.

The memory 102 is... MSG(underscore)1, in accordance with the invention, allows the performance of the message passing process to be linked to the relative complexity of the message that is to be transferred... information, such as whether the message is RPC or IPC, whether it is from the server side or the client side, is always present and can be found in the primary... The MCS is used by the message passing library 220 to translate ports and to process by-reference parameters to emulate local procedure call semantics. The TCS and the information it... is to minimize the amount of information that must be recopied in the message passing process.

Values for the destination task's 210" port name, the name of the message, and... transmission from the sending task 210" of Figure 7C, to the IPC subsystem 122. The process is carried out in a data processor such as host 100, that

includes a memory 102 in which there is resident the... associated with the first task 210, for executing the instructions from the program in a processor 110. This is also an object that is provided by the program at compile time... at runtime.

Step 768 then executes a first instruction in the thread 248 with the processor 110, to load a data value 720 into the send data buffer 752 and to... 770 then executes the procedure call mk(underscore)msg in the thread 248 with the processor 110, to make the header template 740 available to the transmission control module 700 of... of the destination task 210". If the message cannot be transmitted, for example because the processor 110 is suffering resource exhaustion, or because a time out has expired, or because of insufficient port rights, then processor 110's time is not wasted in an abortive transfer of the data portion 720... 7D shows the substitute transmission control information 708, as it appears after it has been processed by the interprocess communications subsystem 112. The substitute transmission control information 708 includes the message... module 700, to be more efficient and faster.

The client task 210 and/or the server task 210' can be part of an application program 180, an operating system personality program... still other tasks concurrently running on different host multiprocessor systems 100', as in the distributed processing network shown in Fig. 8. Each communication from one such task to another can avail... 8 shows a functional block diagram of two host multiprocessor systems running in a distributed processing arrangement, with the IPC subsystem 122 and the transmission control module 700 on each host processor managing interprocess communications between tasks with the exchange of messages between the two hosts over... to the task 211' in host 100' in the manner described above, for a distributed processing application.

The microkernel also includes the capability engine module 300 that manages capabilities or rights... United States Patent Application by James M. Magee, et al. entitled "CAPABILITY ENGINE METHOD AND APPARATUS FOR A MICROKERNEL DATA PROCESSING SYSTEM". Also see the copending United States Patent Application by James M. Magee, et al. entitled "TEMPORARY DATA METHOD AND APPARATUS FOR A MICROKERNEL DATA PROCESSING SYSTEM". Also see the copending United States Patent Application by James M. Magee, et al. entitled "MESSAGE CONTROL STRUCTURE REGISTRATION METHOD AND APPARATUS FOR A MICROKERNEL DATA PROCESSING SYSTEM". Also see the copending United States Patent Application by James M. Magee, et al. entitled "ANONYMOUS REPLY PORT METHOD AND APPARATUS FOR A MICROKERNEL DATA PROCESSING SYSTEM".

The invention applies to uniprocessors, shared memory multiprocessors, and multiple computers in a distributed processor system. Figure 8 shows a functional block diagram of two host multiprocessor systems 100 and 100' running in a distributed processing arrangement, with the IPC subsystem 122 and the transmission control module 700 on each host processor managing interprocess communications between tasks with the exchange of messages between the two hosts over... 250.

In Figure 8, the thread 248' of the host 100 sends a request for processing to be executed in the I/O adapter processor 108. The instructions sent by thread 248' for execution can include those necessary for the formulation of a message to be sent from I/O processor 108 to I/O processor 108' of the host 100'. Such a message can be a message for a server task 211 or 211' sent from the client Task(A) 210 to Task (B) 210... discussed above. The message is sent over the communications link 250 to the I/O processor 108'. There, the thread 249 associated with task 211, executes in I/O processor 108' and transfers the information from the message to the task 211. Another IPC transfer... 100'. The thread 249' belonging to task 211' in host 100', executes instructions in the processor 112' of host 100', and can operate on the information in the message it receives... facilitating interprocess communications either within its own memory 102 of residence or alternately with a processor 112' having separate memory 102'.

Section 1: Subsystem Level Interaction

The IPC 122 subsystems relationship... to skip the explicit message creation step altogether when the sender and receiver are both synchronized. These variants represent internal optimizations of the IPC 122 library which are

transparent at the component level. The conditions under which synchronization is experienced and the opportunities created by it are explored later in the paper but in general synchronization is present in all RPC cases, most send/receive IPC message type cases, synchronous send...in which data contained in the by-reference buffers to be used, i.e., stateless

servers have no need for data associated with a call after the reply message has been... ..makes it possible to re-use the associated allocated address space, hence the existence of Server Temporary. In another example, receivers which are acting as intermediaries or proxies may not need to access a data region and therefore have no need...
...400:

In RPC (remote procedure call) transfers, many send/receive pair IPC message type (Inter Process Communication) calls and some apparently asynchronous IPC message type calls, the receiver knows ahead of... ..may start without regard to boundary and that it will be concatenated together with other server temporary parameters in a region of memory provided by the receiver on a per instance... ..earlier CMU message passing system and gives a method for call specific (or more usually server specific in the case of demultiplexing servers) disposition of individual parameters.

2.3 Shared Data

The shared data class requires specific setup...of the performance implications but if these are found acceptable a non-local client or server is possible. Further, since we have established a formal language for describing the portion of... ..utilized in a fashion which is transparent to the two application level parties.

2.4 Server Allocated Resources

As its name implies, this buffer sub-class is specific to RPC. In... ..the called procedure). For cases in which the buffer is to be provided by the server it is necessary to suppress buffer allocation by the IPC 122 subsystem. To enable the... ..will want to suppress IPC message type level buffer allocation through the use of the server(underscore)allocated option. Even if we were willing to accept the server side always expecting a buffer and having the library routine for the local call create this buffer, there is still a performance related reason for suppression. The server may already have a copy of the data the client wishes to see. Full support of the server(underscore)allocate option means that the server is allowed to set the client sent parameter to point directly at this data, this is obviously the method of choice in a local interaction. If we always required the server to accept an incoming buffer, the local case would suffer. The intermediate library routine would be forced to allocate a buffer, and the server would have to copy data from its permanent source into this buffer. A similar scenario... ..and though it slows down the transaction, it is less performance sensitive.

2.5 Sender (Server) Deallocate

The sender deallocate buffer subclass is present in IPC 122 and on the client... ..deallocate a buffer pointed to by one of the calling parameters. Without the availability of server(underscore)dealloc, support of this behaviour in the remote case would require explicit buffer deallocation... ..the send before returning to the application. RPC also supports an analogous option on the server side dubbed server(underscore)dealloc. Server(underscore)dealloc can be used on buffers associated with data the server is returning to the client, with buffers the server is receiving data on or buffers which serve both functions. In the server send case serv(underscore)dealloc behaviour is the mirror of send dealloc. In the client send case, the server(underscore)dealloc function appears to operate like a server associated with a server(underscore)dealloc follows the rules of permanent buffers. This makes it easier to manipulate on subsequent calls within the server. Further, the buffer which is deallocated when the server makes its reply is the one associated with the reply data, not necessarily the one...option was set. Non-direct data associated with extended transmission control information is sent as server temporary. This means it will show up in the same buffer the header and extended... ..to be shared between the sender and receiver. Passing of a shared capability to a

server does not require the server to make provisions ahead of time. The server will be able to detect the share setting of the capability and will take whatever... ..interfaces. While it would be straightforward to emulate RPC, more complex IPC message type, passive servers, etc. on top of the capability engine 300's primitive message service, there would be... ..an income send message but this time the capability engine 300's check for waiting servers meets with success. The proper server is targeted for thread handoff through the customizable queue call. The capability engine 300 now... ..the direct transfer of capabilities and ports but also the mapping of a capability (the server requests an income capability be mapped) or the unmapping of one. (the server request an incoming by-reference buffer be received as a capability) Scheduling again takes place... ..engine 300 unless the client is to wait upon an explicit reply port. (If the server is not accepting anonymous reply, not guaranteeing the reply will be returned by the entity now receiving the message.) The scheduler is called if the server is to run with the receiver's scheduling properties.

In example three, of Fig. 12... ..responsibility to guarantee against the arrival of a send while the receiver is in the process of blocking, or to check once more for senders after the block has occurred.

Example... ..message is either cobbled together directly (IPC message type send/receive) or created in a server loop designed to assist the target end point of an RPC. (emulate a local call...within a compressed and carefully thought out structure is significant from an interface definition perspective. Server's can check the match between the message format they expect and the one the... ..of the message control structure sent by the client. When a match is found, the server will be guaranteed that pointers within the message will be pointers, ports will be ports... ..not guarantee semantic meaning to the associated data of course, but it does mean the server is protected against random pointers and random values for port rights. The server is guaranteed of this because the message control structure (and server provided overrides) hold the sole determination criteria for the message parameter format.

What further makes...message type model. A simple message may still require a message control structure if the server wishes to test it for compatible format. This should be a very limited case however. If the server is expecting a particular message or recognizes a group of message id's, a simple... ..behaves no differently than one which simply contains garbage data. The only case where a server might need a message control structure is on messages containing variable simple data format not... ..which accept messages which are not pre-defined. Without the definition of every parameter, the server would not be able to parse the incoming message. There have been efforts to improve... ..message control structure dictated expecting where by-reference and capability disposition has been influenced by server use of the overwrite buffer.

If the caller wishes to know or influence some specific...objects) requested. This implies that by-reference regions associated with transmission control parameters are considered server temporary. As will be detailed later, when a server in an RPC or the target of a 2 way IPC message type calls the... ..prepared to accept all of the transmission control information as outlined above, but also the server temporary data 400. The format of the returned buffer is, header at the top followed by direct optional control information, followed by server temporary fields, including those associated with the transmission control information.

Figure 21 is a diagram... ..receiver message format and identification. Some punt, assuming that the question of message format is settled by the sender and receiver outside of the message passing paradigm. Others pass partially or... ..In an embedded system sending fully trusted messages, it is hardly necessary to burden the processor with generic message parsing. On the other hand, in the general message passing operating system... ..the receiver must verify message format. General message passing also makes use of generic receive servers which parse a message to determine its format. With the separation of message control information... ..passing library 220. This convention is absolutely necessary in the case of asynchronous messages where server input simply may not be available. Although not absolutely necessary in the synchronous cases it... ..is likely to be lower. If the client sent a message and relied on a server message control structure to parse it, some percentage of the time an incorrect message would... ..client message parameters. A non-trusted client would then be sending garbage data to a server. If the client is required to send a message control structure, the server checks the non-trusted client message control

structure, avoiding the receipt of garbage data. (The... ..client. If the client were to send a message to the wrong port in the server message control information paradigm and that message were to unintentionally ...data to unmapping and overwrite, i.e., a client may send a message to a server, expecting that there are two direct parameters. The server believes the first parameter is a by-reference and that further, the associated buffer is... ..consults the sender supplied message control structure to translate all non-direct data parameters. The server, however, is expecting messages of only one format, or in the case of a demultiplexing server, messages whose format is determined by the message id. The server, therefore, does not request the message control structure and acts on its assumptions. Such a server could be damaged by a client either intentionally or unintentionally sending a message of the wrong format.

With the receipt of the Client's message control structure the server

is now free to check the format of the incoming message against expectations. If the server is demultiplexing, the message id is checked first to determine which amongst a set of... ..message data as shown in Fig. 23. This last scenario is most likely when the server is acting as an intermediary for another server. The use of the message passing interface to implement a communications server can make a good example of the power of the message passing library 220. For... ..communication code.

3.4.1 A Fully Defined Send-Receive Compatibility Check

Even if a server and client have fixed on a message format, or in the demultiplexed server case, a series of message id pared formats. The server may not trust the client to do the right thing and send the appropriate message... ..incoming parameters contains no other information. This makes it possible to do binary comparisons of server stored message control templates with the incoming client message control structure. The distillation of message... ..and buffer disposition for both the request and reply. It is very reasonable that a server would want to support clients that chose different local buffer disposition options. As an example let us consider 2 clients which both want to interact with a common server. They both want to send a by-reference field to the server. One wants the buffer removed after the send, the other wishes to retain it. It would be awkward if the server were to reject one of these two clients just because neither of them was trusted... ..in the template and the client derived message control structure before the binary check, the server can service both clients in non-trusted mode.

The example above brings out one other... ..and the client option mask override for RPC message control structures as callable macros, the server is free to fashion any sort of partial check it sees fit. For example, allowing... ..call basis, even on complex messages. The method involves message control structure registration 500.

A server wishing to participate in registration, makes a registration call for the message control structures associated with the server's set of interfaces. The registration call parameters are the message control structure, the associated... ..of the port. A client wishing to send messages via the registration service, contacts the server with a simple call, sending the message control structure; possibly containing a message id, and asking for the associated registration number. The server is free to run what checks it likes, but in practice absolute compatibility is required. Should the server detect for instance a difference in client local buffer disposition and pass back the registration id anyway, the client would be damaged upon the use of that registration id. The server may fail a registration request which does not match exactly or register an additional message control structure for that particular message id. The server would then be responsible for checking both registration numbers for that particular message id, the server template registration number and the client registered on. The server should also keep a copy of the client message control structure on hand to check... ..is still free to attempt non-registered transfer.

The registration of message control structures for servers which persist over long periods is certainly indicated for both trusted and non-trusted client-server pairs. It will be most significant in the non-trusted case, however, since it removes the need to copy the message control structure to the server and do the call by call check for format compatibility. A registered server will work with both registered and non registered senders. Therefore, if a sender is

only... 220 is set up such that the client must request the registration information of the server for two important reasons. First, it reduces the code which must be maintained in the message passing library 220. Second, the server maintains full flexibility in determining who matches registered message formats and who does not. Use... make a wide variety of incoming message formats compatible. It is up to the individual server to sort through this and support the set of formats it sees fit.

3.2... overwrite buffer are 1: Permanent data (note: the permanent, by-reference choices also include the server dealloc cases.) and 2: capabilities. It is possible via the overwrite buffer to request that... exercised.

Use of gather often necessitates the request of send message control information by the server so that the actual size and number of permanent regions and capabilities will be known... with dynamic by-reference regions.

In the case of RPC it is necessary for the server to construct a message buffer for the reply which in the format the client is... an example of overwrite use.

3.2.4.3 Reply Overwrite Control Information:

When a server redirects the placement of data on a by-reference region using the overwrite option, care must be taken to ensure the post, or reply processing is appropriate. An RPC style interface might well have been set up to deallocate a by-reference region using the server-dealloc option. If the server has re-directed by-reference data to a region that it wishes to persist past reply delivery, it must pass back an altered message control structure. Upon detection of the server side reply side control structure, the message passing library 220 scans it for server side buffer disposition overrides. The message the client is expecting back in the case of RPC is of course in the client format. It is up to the server to put together the appropriate message buffer. It might have been possible to send null buffers back on fields upon which the server-dealloc option was set for buffers which were only passing information to the server. This however, was an insufficient answer for the buffers being used to send data both... yet ready to receive the data. Synchronous interfaces on the other hand need never create intermediate capabilities for the by-reference data types, because the sender must pause for a pending reply anyway, the point of synchronization for the client is not the return from the send but the return from the... is available and the transfer can proceed from task space to task space without an intermediate message. It is clear then that the message passing library 220 must also formalize the... associated data to the task space where the remote procedure lies, await the remote procedures processing, return the results to the callers space and finally make all the incidental changes like... client initiating a call will not succeed in starting that call and activating the associated server in a non-restartable way only to find out that a loosely paired reply does the initiation of the RPC has an impact on message checking. A server utilizing the overwrite option may accept a wider range of incoming client messages and may... than one message control structure format grows out of the asymmetric nature of the client/server relationship. The server registers the client's message control structure. If there are two clients which send exactly the same format message but wish to receive the reply data differently, the server must register two message control structures to support them both.

The implications of restriction 1... transfer. The message control structure is kept in the message passing library 220 while the server is active in anticipation of the reply. If the message was complex but was accompanied by a large amount of direct data, the server can avoid sending this data back on the reply by sending an override message control... the by-reference, capability, and other port fields in the message buffer sent from the server and will fill in client buffers, or update the client's double indirect pointers as... are two major models of RPC support with respect to scheduling. The active and passive server models. In the active case the scheduling information associated with the client's request is that of the server thread. In the passive, it is that of the client. In the active model, the server can be observed to directly commit a thread to the receipt of a message on... client then sends a message to this port and blocks waiting for the reply. The server thread returns to non-supervisor mode with the message and proceeds to process it, returning with a reply when processing is complete. In the passive model, the server as owner of a port, prepares a thread body, (it prepares state and a set... level thread). The client does not so much send a message as enter the target server's space with the kind of restrictions associated with a traditional kernel

level service call, i.e., start execution at a target mandated point, process incoming parameters along previously defined lines.

In the case of RPC the assurance that the client will block while the server is working on its behalf is very helpful in supporting elements of the passive model... temporary resources associated with a client thread at kernel level may be borrowed by the server. The thread stack and other temporary zone space are good examples. The client prepares a message for transfer, the server is then allowed to borrow the buffers which hold the results of that preparation. In... of thread migration as separate options. The most important is, of course, scheduling. If the server thread in the active case inherits the client's scheduling characteristics, and the kernel elements... and active models.

In the active model an actual runnable thread is created on the server side. This may or may not be used for other activities, in either case it... client loans its kernel temporary resources and its schedulable entity, effectively its shuttle to the server thread, now effectively a thread body. The client entity, now effectively a thread body is... expensive in the exposed case. This might give a small advantage to the extremely dynamic

server case. Depending on the exact nature of the interface it is also possible in the...be supportable without such direct exposure. See section 4.1.9

4.1.2 Client/Server Juxtaposition:

Client/Server juxtaposition is characterized by the synchronization of the client send and server receive. In the case of RPC, if the server arrives at the receive port before there are any messages to receive, it blocks. If... until the receiver arrives. This in effect guarantees simultaneous access to both the client and server space for the purpose of message transferral. Though client/server juxtaposition can be achieved in some circumstances in asynchronous communications it cannot always be guaranteed... avoid capability translation by temporary conversion to direct data. This seems especially likely for the server(underscore)temporary examples.

Assuring client/server synchronization also reduces the need for kernel level resources and leaves the remaining resource needs more... be easily constructed. Thread A sends a message to thread B. B, however, is busy processing an earlier request (possibly from A). To process this request, B must post messages to several other tasks. Each of these messages requires... multilevel operations to reserve all the necessary storage before beginning but this sort of transaction processing has problems of its own and is, of course, inherently synchronous in nature. Client/Server synchronization can reduce kernel resource requirements to some small number of bytes per thread in the... The first is the issue of verifying the compatibility of a client interface when the server is using override options to alter the disposition of local buffers associated with the call... valid. This problem has been gotten around by collecting all of the bits associated with server buffer disposition into a field. The server may check the incoming message control structure with by the same byte by byte comparison... for the addition of a masking operation before the comparison of parameter flags fields. The server is in full control of the compatibility check, based on the type and scope of... some of the parameters of the incoming message. The second drawback is centred around the server(underscore)dealloc option. Special care will have to be taken when it comes to server(underscore)dealloc, the server may be compelled to check for this option and where it occurs send an override... sub-optimal in the sense that if a client persists in sending messages with a server(underscore)dealloc on a parameter and the server persists in doing overrides in which the server(underscore)dealloc must be overridden. The server must continually send a reply message control structure and the message passing library 220 must on a call by call basis consult it. In the worst scenario, the server would check the incoming message control structure every time, doing a special check for the... not a large disadvantage over a hypothetical non-coalesced notion since in that case the server would have to send a message control structure with every reply. But it does require an extra check at user level on the part of the server and a cross comparison at the message passing library 220 level. This, of course, can be avoided by having the client send a message control structure that does not contain

server dealloc, or through registration. The server can choose to register a message control structure which does not include the server(underscore)dealloc option and return the registration id for this to the client.

4.14... ..routine is capable of supporting local semantics, i.e., if the client was expecting the server to use a particular heap source when allocating a buffer, the proxy might allocate such...case of RPC, rather, it is expected that SVC calls which encounter a shortage of server resource (either active thread or passive thread body notion) will trigger the capability engine 300... ..300 because ports are only accessible through capability calls.

4.1.7 Support For Message Server Spaces, Demultiplexing on Message ID:

It is often the case that a series of functions...and the message control structure must be consulted.)

In the demultiplexing model, the user level server is comprised of a primary procedure which places messages on and recovers messages from a port. This procedure does some general processing, (message too large, restart handling, buffer handling, etc.) and in turn, calls a server side proxy. The proxy called is determined by the message passing id. This server side proxy does the discrete function level specific setup and checking.

The general server loop does get involved in discrete message processing in one place. The message control structures are made available to it through a table... ..is generic. It is just the data involved which is function specific. Further, alterations to server side options are necessarily server wide in scope. It is the general server function that is the most appropriate place for the necessary adjustments to the server side format check. It is also true that the server side stubs tend to be automatically generated. This makes them a less convenient target for... ..flow of execution in a typical message receive/send, as shown in Fig. 28.

- * Primary Server Function receives Message

- * Primary Server Function checks status, i.e., message to large and does appropriate high level handling.

- * If registered, primary server checks index table to relate registration id to message id.

- * If not registered and client not trusted, Primary Server Function uses message id to get message control structure template and check against incoming (obviously requested) sender message control structure.

- * Primary Server Function uses message id as offset into table to get the proper proxy function. Primary Server calls proxy function.

- * Proxy function does any necessary transformations on incoming data. These transformations are... ..callee)

- * Proxy function does function specific cleanup, including any data transformations. Proxy function returns.

- * Primary Server Function re-works header fields, it is not allowed to increase the size of the header unless another buffer is used. (There may be server temporary data 400 to be sent on the reply below the header in the receive buffer.) The primary server optionally includes a reply message control structure (rare) and reworks the message returned by the... ..be supported directly in the product, the application writer may be left to customize the server loop and data structures by hand.)

- * The Primary Server Function calls the message passing library 220 with a send/rcv. The supplied header is... ..or one of the oversize options may be encountered.

4.1.7.1 Dynamic Message Server Spaces:

The support of dynamic linking and co-residency is very powerful. It allows the... ..perform possibly as a local procedure call without any additional overhead, effectively bypassing client proxy, server, and server proxy routines. When the function call is aware of message passing, it will still be... ..reduced when contrasted with a remote call.

Co-residency also supports the remote setup of servers. To support this, co-residency must go beyond simple download and link functionality. In the case of a pre-existing server, download and link from an external source could be used to alter one or more of the server proxies and their endpoint routines. In order to do this, however, the remote entity would... ..communication between the target and the task attempting remote download. The caller cannot start a server, or add a new message id to an existing server without some method outside of the defined notion of co-residency.

In order to support these notions in a simple straight forward manner, we need support for a dynamic server model. A task wishing to make itself available as a dynamic server must create and export a port which makes the series of server creation, manipulation, and shutdown routines available. A server for servers. This server/server exports the calls presented by the server library. The default server loop is not just a shared library routine. The `server(underscore)create` call creates a threadless instance of a server and returns a handle. This handle is used by subsequent calls to change optional aspects of the server instance, add or delete server threads, associate proxies and by consequence their endpoints, add or remove receive buffers, or shutdown and clean up the server instance. After using ...utilities to download specified code into a target task, the remote caller would send a `server(underscore)create` message to the server/server port and receive a handle back on the reply. The caller may have supplied a... ..caller has an additional call which is not one of the calls exported by the server package. An extra call is needed to create a thread and then direct that thread to associate itself with the target server instance. In the passive model, it is possible to simply provide the thread body resources to the receiver, but in the active model, the server acquires threads via a call by the target thread. There is an advantage to having the routine built in this way. The target server task is free to adjust post processing or customize thread state or resource for its specific needs. Because of the notion of server instance, a server persists even if its threads exit the server. In this way, exceptional conditions can cause a thread to return from its `run(underscore)server` call. The task is then able to customize exceptional processing. The thread can then be returned to the server loop. If the exception the thread is returned on is a simple `return(underscore)server(underscore)thread`, the thread is free to re-associate itself with the server, run some other unrelated task or self-terminate.

4.1.8 Anonymous Reply Support:

In... ..The client is simply blocked waiting for the completion of the remote procedure call, the

server or at least a thread of the server is dedicated for the duration of the call to completing the remote procedure call and... ..do the wait and to map a send or `send(underscore)once` right into the server's space. There are cases, however, where for throughput or transmission control reasons more flexibility... ..of this flexibility, in some circumstances, an explicit reply port is required on either the server or client side in order to keep track of the reply target. Though rare, the client may wish to declare an explicit reply port in order to allow for intermediate message delivery. The proxy routine would in this case be capable of receiving these intermediate messages, processing them and then re-establishing the wait for reply by doing a receive on the... ..block on send (or request), a block on receive (or reply) and possibly out of server processing via an out of band abort signal to the application server routine. If the client side proxy is set up to handle it, the send side abort with signal is straightforward. The client awakes with an `abort(underscore)notify` signal, processes it and if it wants, restarts the RPC. If the server is already processing the request, however, the client is waiting on a reply, in order to receive a... ..the `abort(underscore)notify` state and include it with the reply coming back from the server.

In order to avoid an explicit reply port on the server side, the server must be able to guarantee that the thread sending back the reply will be the... ..way, the client awaiting a reply can be registered in a structure associated with the server thread structure. The server may not be able to guarantee this as it may be subject to a user... ..optimization requires that a client decision with respect to reply port be hidden from the server and visa versa. The

message passing library 220 achieves this with the algorithm shown in...at the point that the message passing system has both the client send and the server receive in juxtaposition. This is, of course, always achieved in RPC.

Case 1, of course... ..4 because it is not necessary to create and place a port right in the server's space. Case 3 may perform nominally better than case 4 because the anonymous port... ..the state and setup of the normal port types.

Upon return from the request, the server thread's data structure is checked for an outstanding reply. This will be present in... ..client is blocked on a port, it is removed. If it is not blocked, the server is made to wait on the port as a sender. When the client is available, the reply is delivered, the client thread returns and the server resources or thread are free to take on another message.

If the server's thread structure does not point at the client, there must be an explicit port in the remote port field of the server message call or an error is returned to the server. The client thread is retrieved from this port if it is there and the transfer proceeds. If it is not on the port, the server thread block awaiting it.

4.1.9 ABORT Support:

Abort support is a complex issue... ..guarantee a quick restartable return from a kernel call so that a user level signal processing routine can be called and returned from. If the thread is running at user level... ..wait that a thread is capable of exercising, a kernel based one and an external server based one. In the case where the caller of abort is not worried about thread restartability, the only important considerations in waking up a waiting thread are the server or kernel resources and state. The thread may be returned to user level with a thread(underscore)aborted declaration at any time so long as the server/kernel are not left in an undefined state or with an orphaned resource. In the... ..of kernel resource recovery is beyond the scope of a paper on message passing. The server case, however, directly involves the message passing system.

In RPC, the abort function may find... ..request, abort is simple and restartable, return the thread with request(underscore)aborted status. The server was not yet aware of the request and no recovery action at all is required... ..In the case of Thread(underscore)abort, an attempt may be made to stop the server as soon as possible rather than letting it complete a now useless bit of work. The first attempt to abort the server is made via the port. A field of the port structure points to an abort(underscore)notify function. If the server wishes to support early termination of work for an aborted client, it may choose this... ..such that when the reply is sent back, the message will be destroyed and the server reply apparatus liberated. If the port is destroyed first, the server will simply encounter a dead name for the reply port and may act to destroy... ..of the receive port has not been filled in, it checks to see if the server requested anonymous reply port. If it did, the server has guaranteed that there is an unbreakable link between a specific server thread and the request. In the server anonymous reply case, the message passing library 220 executes a thread(underscore)abort(underscore)safely on the server thread and sends a signal indicating that the message being processed is no longer important. The anonymous reply port, if present, is destroyed. If the client... ..the reply port is set such that the reply message will be destroyed and the server reply apparatus liberated as if the reply was sent.

The client will return from its message passing... ..retry can be achieved if the client checkpoints its data before attempting message passing. The server state is important only if the server maintains state between invocations. In this case, the designer must insure that the server receives notification of client aborts and takes appropriate action.

From a real-time perspective, there is a danger to proper scheduling of resources in the case where the server acquired the scheduling properties of the client. From a scheduling standpoint, this is effectively the server passive model where the client entity runs in server space. After experience of an abort. The client thread is effectively cloned with one running temporarily in the server and one running in the client. If the priority of the client is high enough, the server thread (in the abort/signal scenario) might run to completion before encountering the signal to terminate. In the server explicit reply case when there is no abort notify port, there is no attempt to notify the server of a client abort.

It is only in the case of the abort notify port that the server, through the system designer, can ensure timely delivery of the client abort notification. If the active thread on the server abort notify port is given high priority or if the passive scheduling parameter assigned by... ..is of high priority, it will be scheduled before and may preempt the client message processing. The server may then set user level state to communicate with the client message processing thread that it must terminate early.

4.1.9.2 Thread(underscore)Abort(underscore)Safely... ..a client experiences a

Thread(underscore)abort(underscore)safely. In the case of the active server, the server thread is free to pick up this request even during an ongoing Thread(underscore)abort(underscore)safely. In the case of the passive server model, unless otherwise instructed (see real time considerations below) a shuttle is cloned and the server processes the request. Another way to look at it is that when Thread(underscore)abort(underscore)safely 1.9.1. The server resources are guaranteed not to become permanently pinned on a non-operating reply port.

There... ..respect to Thread(underscore)abort(underscore)safely. It could be argued that the asynchronous signalling process is a schedulable event in its own right. The exception message port would carry with... ..a need to adjust the scheduling information of the RPC at least in the passive server case. If the RPC has not had its request considered, the scheduling info can be altered to reflect suspend. This, of course, may effect the order of queued requests processed. If the request is already underway and the client abort notify port is active, the message service can send a message to the server stating that the request should be suspended. If the client notify port is not active and the server is using anonymous reply, the server thread may be suspended. It is assumed that the first non-intervention approach will win... ..receiving a thread(underscore)abort(underscore)notify(underscore)send message. There are no artifacts, the server was never aware of the request, and the client is free to retry the RPC... ..status of thread(underscore)abort(underscore)notify(underscore)receive. The client is then free to process the notification and do a receive on the reply port to continue the RPC. Thread... ..system designer may feel compelled to take some action to effect the scheduling of the server request handling of the associated RPC, especially in the case of the passive server model. The case differs from Thread(underscore)abort(underscore)safely in that the synchronous code... ..unless either the sender is using the anonymous reply port option, in which case the server thread may be signalled, or the client abort notify port is active so that a client abort message (of notify flavour) can be sent to the server.

4.1.10 Shared memory support 800:

Shared memory regions may be established through the... ..two means. 1: The explicit shared memory by-reference parameter with matching overwrite on the server side to establish the shared region; or 2: The passing of a share capability. Though... ..in detail in this section. Passage of a share capability is less constraining to a server. The server is free to send the capability on to another task and if the write was a send right instead of a send once, the server may share the share region with others. This will create a common buffer among multiple parties. The

server does this by doing a copy send of the capability in a message passed to...the sending of a message, broadcast is not defined for RPC and only the target server will receive an overt signal that the message buffer has been filled or freed. Capabilities... ..therefore, protected from broadcast in the explicit shared memory setup case.

The vast majority of servers are stateless, that is they do not hold information acquired specifically from the data associated... ..stochastic information concerning frequency of usage and resource requirements.) Because of the preponderance of stateless servers, it is expected that cases of shared memory usage in RPC will strongly favour client... ..and location of the shared region within their own space. In this way, if a server did not find that a particular client was trusted enough, it could decline to accept... ..up to be shared. This is important because it is often true in the active server case that a server cannot trust the backing pager of a region of shared memory offered by a client. The associated pager might be inactive and cause the server to hang on a memory fault. In this case, it must be the server who offers a region. The region, of course, will be backed by a pager that the server trusts.

Because of real time considerations, client pager issues notwithstanding, it is likely that the client directed share will be the method of choice in the passive server model. The passive model asserts that the client thread of execution enters the server space. In this case, it is the server pager that the client might not trust. Should the thread associated with the client request... unmapped memory.

In Fig. 31, Message Passing library 220 does not see overwrite buffer on server side; it checks to see if client region is actually shared by server by checking task specific state info setup on share region initialization. If region shared, it... it. Either party may initiate the call, i.e., if a client existed in the server task space which called a server in the client's space, the message passing library 220 would support it. Further, all... oxymoron. There are, however, conditions which require support of one-way messages, especially on the server side. In the active server paradigm, the server thread usually performs a reply with a subsequent receive for the next message. This is... state condition, but how do you start it? The answer is that you start the server with a one way receive. The thread resource is put to sleep awaiting the arrival of the first message.

Server's may find it convenient to process subsequent messages while a particular message is blocked, waiting for a resource. Further, they may... a reply to get the next message. When the blocked message is finally reactivated, the server finds that it must now do a one-way send to institute the reply.

There...order to show the flexibility of the RPC interface.

The RPC interface does not collect server interfaces together or create the proxy function tables and message control structure tables. It is... data types. The only by-reference data type not supported by IPC message types is Server(underscore)Allocate which exists in RPC because of the semantic link between request and reply... IPC message pass.

The header, message control structure, and message are shown as a contiguous unit because it is expected that in a large number of IPC message types cases, the... may be used for outgoing optional transmission information because even in non-proxy uses, IPC servers may wish to make receipt of such transmission information because even in non-proxy uses, IPC message type servers may wish to make receipt of such transmission information optional. (If the transmission information section...as the sender intended.

* Immediate update of non-local share regions: Message presentation represents the synchronization point from the sender's perspective. This may result in slightly different behaviour between the... will not matter because the behaviour is undefined unless the sender refrains from altering the buffer until it encounters a separate synchronization event. This synchronization event will not occur until after the receiver has received the message and processed the data in the shared region.

The receiver may attempt a receive before or after... into the receiver's space, an the message buffer, message control structure if requested, and server temporary data are moved into the receive buffer, as shown in Fig. 33.

As in... way message and would be altogether unremarkable from an implementation perspective except that the enforced synchronization makes it possible to avoid explicit message creation. This allows synchronous one-way IPC message... a substantial advantage in system level resource management.

In both the synchronous and asynchronous cases, server(underscore)temporary usage will be impacted by application level convention between senders and receivers, knowing...way)

2. Party 1 does a send then a receive, Party 2 does a receive, processes the data and then does a send. (Supported by synchronous and asynchronous 2-way. As in the RPC case the second party starts out the process with a stand-alone receive.)

3. Party 1 does a send then a receive, Party... send is done by a different thread. (Supported by asynchronous and synchronous 2-way, but server cannot use anonymous reply port option)

4. Same as three but the incoming and outgoing... ..2-way would wait on the reply, the local caller would miss the opportunity to process the remote data while the data it sent was being processed.)

It should be stressed that two-way IPC message type support is a performance optimization... ..the local or reply port disposition information field in the header.

4.2.2.1 Synchronized Two-Way IPC message type:

Synchronized two-way send enjoys performance advantage because of the separation of the system level semantic linkage of request and reply as in RPC from the issue of send/receive synchronization. The IPC message type linkage of synchronous two-way send dictates that the message control... ..that takes the receive will do the linked (only in an execution sense) reply. Thus synchronized two-way IPC message type gains another performance advantage, enjoyed by RPC.

Synchronized two-way IPC message type differs from RPC in that the two parties of the... ..to party 1. Whereas the RPC transaction requires that message buffer data returned from the server be passed back as indirect data, the IPC message type transaction is free to pass...thread can still pick up its data, return from the IPC message type and continue processing.

4.2.3 Message Control Information Specifics

The format of the message control structure in... ..control structure parameter descriptors. It is conceivable that an application writer might want a particular server to accept both IPC message type and RPC messages. The port could be some sort of intermediate collection point, it is for this reason that the message passing library 220 avoids rejecting... ..control structures also makes the reply override found in RPC unnecessary. The only reason a server sends a message control structure to the message passing library 220 is to do an override of server side buffer disposition options (specifically server(underscore)dealloc, please see section 4.1.3). This situation does not exist in IPC... ..determined by the message control structure in a fully symmetric reflection of the initial send process. The main reason for override use in RPC is the initial customization of server side buffer disposition on the request by the overwrite structure. The overwrite structure contains server specific requests to alter the placement and class of incoming by-reference variables as capabilities... ..a registration value. (Please see section 3.2.4.3 for a description of the process of registration.)

As may be recalled from the section on general registration, the client (or sender) must send the server (receiver) a copy of the message control structure and request the associated registration handle. The server is free to: 1) deny the request; 2) comply with it after matching the entry...to the OUT parameters. There will be no IN/OUT and IN parameters. Further, the Server(underscore)Allocate class does not exist. The reason for the Server(underscore)Allocate class is that a parameter with a buffer description exists into which the server is to place data for the client, and the writer of the interface does not want the message passing library 220 to provide the server with the buffer on the request. This is done for performance and to increase the... ..a direct outgrowth of the knowledge of the reply at the time of the request.

Server(underscore)Deallocate survives as a data type only in the form of a Sender(underscore... ..can be characterized as data passed by-reference which nonetheless will not persist beyond the processing of the data associated with the message, or at least will not be required past... ..following message. In RPC this sub-class of by-reference support is referred to as Server(underscore)Temporary. The data in this subclass might be used in a reply to the... ..be reused.

IPC message type non-persistent message data shares the same mechanism as RPC Server(underscore)Temporary and its circumstances and ultimate treatment are similar. Like Server(underscore)Temporary, IPC message type non-persistent memory is placed in the receive buffer after the header, message control structure (if requested), and message body. Beyond the restriction that Server(underscore)Temporary or non-persistent memory data must follow the three structures enumerated above, they may appear in any order in the remaining space. The main different

between Server(underscore)Temporary and non persistent data types is when the data becomes stale. In the case of Server(underscore)Temporary, the data may be used on the reply. In the non-persistent case, the data loses its importance as soon as the associated send has been

processed. The IPC message type is then free to do what it likes with the buffer as soon as the send is processed. The RPC must keep the receive buffer until after the reply has been sent. (This... ..physical memory associated with the parameter may actually be shared between the client and the server. In this way shared memory remains transparent to the application at least at the interface... ..the shared region. This may be done in a separate service, however, allowing the endpoint server and client code to remain free of share specific protocol considerations. If the application ...On initialization, the client sends a by-reference pointer with the share option on. The server sets up an overwrite in which it indicates that it will accept a by-reference... ..members do not share physical memory they are updated. Only one entity, however, the target server of the call, will be sent a message.

IPC message type semantics do not restrict...arrive earlier than messages from that source are read by the remote party. Use of memory update only requires data level synchronization flags and buffer monitoring, or some other external form of synchronization.

Once a shared multi-way has been set up and one of its existing members does not share a common memory resource, extending membership will require synchronization of the membership lists. When adding a new member, the message passing library 220 traverses... ..capabilities on the receive side are identical to those available in RPC. In the RPC server receive case and in IPC message type, the receiver communicates its wishes through the overwrite... ..must prioritize them properly based on their scheduling information.

The choice of running the remote process using the scheduling information of a remote thread or the scheduling priority of the message... ..of the send or elsewhere, the act of queuing may have to wait. Thus, active servers running at high priority may block a medium priority task in favour of carrying out the processing of a request for a lower priority entity. This is an issue of personality level system design, however, and does not impact directly the queue procedure. Unfortunately, in the passive server case, if the incoming request has a higher priority than tasks currently running in the... ..the target, or it may call an associated high priority call back in the target server, allowing the server level code in the target to make specific adjustments.

IPC message type servers are more likely to be active in nature (running with their own scheduling policy) but passive servers receiving asynchronous messages are not unreasonable. To the calling application, the scheduling characteristics in a...possible to provide shuttle like passing of kernel level thread resources between the client and server without it.

The reason that the bulk of thread migration mechanism can be hidden from... ..level interface in the RPC case is the link established between the client and the server and the fact that it lasts for the duration of a full send/receive transaction... ..abort, (please see section 4.1.9 for details) the client remains suspended while the server acts on its request. This allows the passing of kernel level thread resources to the server without visible artifact. The server will process the request and the client will get back its shuttle when the reply message is sent. The passive model where scheduling information is transferred from the client to server for the duration of the request processing behaves in a simple and predictable fashion, following the model of a travelling execution entity moving from space to space in its effort to process its data. The only IPC message type transfer style which behaves as RPC is the... ..party's side and the local party continued with its shuttle to pick up and process the message on its receive port. In this model new shuttles spring into existence based largely on the stochastic elements surrounding message creation and processing within the two parties and their scheduling priorities, and then wink out when a send... ..is passive and scheduling information is being transferred from the sender, the threads of the server may become dominated with high priority scheduling information. This along with simultaneous running of the... ..behave identically when it comes to resource and scheduling issues of migration. The delay and synchronization event experienced in one-

way synchronous IPC message type only affects message creation. It is still the aim of the call to let message processing and the callers subsequent processing continue simultaneously.

One-way IPC message type, in a non-migratory model where scheduling information... ..message looks very much like a thread fork. For the duration of the remote message processing, a remote thread and the local caller will be running at the scheduling priority of... ..the shuttle of the sender. The act of a receive frees the shuttles of its patron and the caller receives the shuttle of the sender when the message arrives. In this... ..chosen, scheduling information is passed from the sender to the receiver. The receiver will thus process the incoming message with the scheduling attributes associated with that message. This works well for... ..thread wishes to do work not associated with a message it has just received and processed, it must do so using the message related scheduling properties or explicitly alter the scheduling... ..in order to satisfy the asynchronous property and allow for the subsequent opportunity to simultaneously process both the post message code in the sender and the message processing code in the receiver. The creation of these messages is not only expensive in processing terms, it also leads to utilization of memory and kernel address space resources which are... ..When a message is created, it has been determined that the sender wishes to continue processing and that a "snapshot" of the data must be taken along with port right movement... ..faster than might otherwise be expected when a system goes from a running state where servers are idle most of the time to one where servers are working a backlog of messages.

4.2.9 Support for Receiver Space Demultiplexed on... ..the message passing library 220, the message passing library 220 is still aware of demultiplexing servers to the extent that a message ID field is supplied in the header. As with... ..the sender and receiver.

4.2.9.1 Support for Dynamic Message Receiver Spaces:

In servers operating on multiple enumerated message formats, it is often nice to customize the handling of... ..has been designed to avoid conflict with message ID customization. Registration is done by the server or receiver and is indexed through a registration ID. This will allow a receiver to...

Claims: ...interprocess communication in a microkernel architecture, the system comprising: a memory means in a data processing system, for storing data and programmed instructions; a data bus means coupled to said memory means in said data processing system, for transferring signals; a processor means coupled to said memory means with said data bus means, for executing said programmed... ..buffer for storing message data information, and having a first thread executing instructions in said processor means, for forming a first message to send to a destination port; a second task... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means; and a transmission control means in said interprocess communications means, for interpreting said message... ..for interprocessor communication in a shared memory multiprocessor, comprising: a memory means in a data processing system, for storing data and programmed instructions; a data bus means coupled to said memory means in said data processing system, for transferring signals; a first processor means coupled to said memory means with said data bus means, for executing said programmed instructions; a second processor means coupled to said memory means with said data bus means, for executing said programmed... ..buffer for storing message data information, and having a first thread executing instructions in said processor means, for forming a first message to send to a destination port; a second task... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means; and a transmission control means in said interprocess communications means, for interpreting said message... ..data available to said second task.

3. A system for interprocessor communication in a distributed processor system, comprising:

a memory means in a first host system of a distributed processor system, for storing data and programmed instructions;

a data bus means coupled to said memory means in said data processing system, for transferring signals;

a first processor means coupled to said memory means with said data bus means, for executing said programmed...
...for storing message data information, and having a first thread executing instructions in said first processor means,
for forming a first message to send to a destination port;

a second task... ...attributes defining said destination port, and having a second thread executing instructions in said
first processor means;

a transmission control means in said interprocess communications means, for interpreting said message control...
...making said message data available to said second task;

a communications link coupling said first processor in said first host system to a second host system of said
distributed processor system;

a second processor means in said second host system, coupled to said first processor means over said
communications link;

said first task providing said message to said communications link, for sending said message to said second
processor means.

4. A system for interprocess communication in a microkernel architecture, comprising:

a memory means in a data processing system, for storing data and programmed instructions;

an interprocess communications means in said memory means, for coordinating message passing between tasks in
said memory means;

a processor means coupled to said memory means, for executing said programmed instructions;

a first task in... ...buffer for storing message data information, and having a first thread executing instructions in said
processor means, for forming a first message to send to a destination port;

a second task... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means; and

a transmission control means in said interprocess communications means, for interpreting said message... ..to said second task.

5. A system for interprocess communications in a microkernel architecture data processing system, comprising:

a memory in the data processing system, for storing information;

a program for a first task in said memory , that includes for the first task, that includes a pointer to the transmission control buffer;

a processor means associated with a thread of the first task, for executing the instructions from the program;

said processor means executing a first instruction in the thread, to load a data value into the send data buffer and to load a control value into the transmission control buffer;

said processor means executing the procedure call in the thread, to make the header template available to... ..by the header template.

6. A system for interprocess communications in a microkernel architecture data processing system, comprising:

a memory means in the data processing system, for storing information;

a program for a first task in said memory means , that... ..transmission control buffer;

a data bus means coupled to said memory means in the data processing system, for transferring signals;

a processor means coupled to said memory means with said data bus means and associated with a thread of said first task, for executing the instructions from said program;

said processor means executing a first instruction in the thread, to load a data value into said... ..control buffer;

an interprocess communications subsystem in said memory means, for managing message transfers;

said processor means executing the procedure call with the thread, to make said header template available to...
...personality-neutral services program.

10. A system as claimed in claim 4 comprising:

a second processor means coupled to said memory means, for executing said programmed instructions;

a third thread in... ..associated with said second task, for providing said programmed instructions for execution in said second processor means.

11. A system as claimed in claim 4, comprising:

said memory means and said processor means being in a first host system of a distributed processor system;

a communications link, for coupling said processor means in said first host system to a second host system of said distributed processor system;

a second processor means in said second host system, coupled to said processor means in said first host system over said communications link, for exchanging said message over... ..interprocess communication in a microkernel architecture, comprising:

storing in a memory means in a data processing system, data and programmed instructions;

executing in a processor means coupled to said memory means, said programmed instructions;

coordinating in an interprocess communications means... ..buffer for storing message data information, and having a first thread executing instructions in said processor means, for forming a first message to send to a destination port;

storing a second... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means; and

interpreting with a transmission control means in said interprocess communications means, said message... ..to said second task.

14. A method for interprocess communications in a microkernel architecture data processing system, the method comprising the steps of:

loading a program for a first task into... ..memory, associated with the first task, for executing the instructions from the program in a processor;

executing a first instruction in the thread with a processor, to load a data value into the send data buffer and to load a control value into the transmission control buffer;

executing the procedure call in the thread with the processor, to make the header template available to an interprocess communications subsystem; and

establishing the transmission...another task in a microkernel architecture, comprising:

storing in a memory means in a data processing system, data and programmed instructions;

executing in a processor means coupled to said memory means, said programmed instructions;

coordinating in an interprocess communications means... ..buffer for storing message data information, and having a first thread executing instructions in said processor means, for forming a first message to send to a destination port;

storing a second... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means; and

interpreting with a transmission control means in said interprocess communications means, said message... ..another task in a microkernel architecture, comprising:

storing in a memory means in a data processing system, data and programmed instructions;

executing in a processor means coupled to said memory means, said programmed instructions;

coordinating in an interprocess communications means... ..buffer for storing message data information, and having a first thread executing instructions in said processor means, for forming a first message to send to a destination port;

storing a second... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means; and

interpreting with a transmission control means in said interprocess communications means, said message... ..another task in a microkernel architecture, comprising:

storing in a memory means in a data processing system, data and programmed instructions;

executing in a processor means coupled to said memory means, said programmed instructions;

coordinating in an interprocess communications means... ..buffer for storing message data information, and having a first thread executing instructions in said processor means, for forming a first message to send to a destination port;

storing a second... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means; and

interpreting with a transmission control means in said interprocess communications means, said message... ..communication with another task in a microkernel architecture, comprising:

a memory means in a data processing system, for storing data and programmed instructions;

an application program means in said memory means, for providing application program instructions to be executed;

a processor means coupled to said memory means, for executing said programmed instructions;

a microkernel means in... ..buffer for storing message data information, and having a first thread executing instructions in said processor means, for forming a first message to send to a destination port;

a second task... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means; and

a transmission control means in said interprocess communications means, for interpreting said message... ..communication with another task in a microkernel architecture, comprising:

a memory means in a data processing system, for storing data and programmed instructions;

an operating system personality program means in said memory means, for providing operating system personality program instructions to be executed;

a processor means coupled to said memory means, for executing said programmed instructions;

a microkernel means in... ..buffer for storing message data information, and having a first thread executing instructions in said processor means, for forming a first message to send to a destination port;

a second task... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means; and

a transmission control means in said interprocess communications means, for interpreting said message...communication with another task in a microkernel architecture, comprising:

a memory means in a data processing system, for storing data and programmed instructions;

a personality-neutral services program means in said memory means, for providing personality-neutral services program instructions to be executed;

a processor means coupled to said memory means, for executing said programmed instructions;

a microkernel means in... ..buffer for storing message data information, and having a first thread executing instructions in said processor means, for forming a first message to send to a destination port;

a second task... ..of attributes defining said destination port, and having a second thread executing instructions in said processor means; and

a transmission control means in said interprocess communications means, for interpreting said message... ..26. A system - 'as claimed in any of claims 23 to 25', comprising:

a second processor means coupled to said memory means, for executing said programmed instructions;

a third thread in... ..associated with said second task, for providing said programmed instructions for execution in said second processor means.

27. A system - 'as claimed in any of claims 23 to 25', comprising:

said memory means and said processor means being in a first host system of a distributed processor system;

a communications link, for coupling said
processor means in said first host system to a second host system of said distributed

processor system;

a second processor means in said second host system, coupled to said processor means in said first host system over
said communications link, for exchanging said message over...

?